

The Efficacy of Self-Supervised Speech Models as Audio Representations

Tung-Yu Wu*

*Department of Electrical Engineering
National Taiwan University*

B08901133@NTU.EDU.TW

Tsu-Yuan Hsu*

*Department of Computer Science and Information Engineering
National Taiwan University*

B08201047@NTU.EDU.TW

Chen-An Li*

*Department of Computer Science and Information Engineering
National Taiwan University*

B08902123@NTU.EDU.TW

Tzu-Han Lin*

*Department of Computer Science and Information Engineering
National Taiwan University*

B08902126@NTU.EDU.TW

Hung-yi Lee

*Department of Electrical Engineering
National Taiwan University*

HUNGYILEE@NTU.EDU.TW

Editors: Joseph Turian, Björn W. Schuller, Dorien Herremans, Katrin Kirchhoff, Paola Garcia Perera, Philippe Esling

Abstract

Self-supervised learning (SSL) speech models, which can serve as powerful upstream models to extract meaningful speech representations, have achieved unprecedented success in speech representation learning. However, their effectiveness on non-speech datasets is relatively less explored. In this work, we propose an ensemble framework, with a combination of ensemble techniques, to fuse SSL speech models’ embeddings. Extensive experiments on speech and non-speech audio datasets are conducted to investigate the representation abilities of our ensemble method and its single constituent model. Ablation studies are carried out to evaluate the performances of different ensemble techniques, such as feature averaging and concatenation. All experiments are conducted during NeurIPS 2021 HEAR Challenge as a standard evaluation pipeline provided by competition officials. Results demonstrate SSL speech models’ strong abilities on various non-speech tasks, while we also note that they fail to deal with fine-grained music tasks, such as pitch classification and note onset detection. In addition, feature ensemble is shown to have great potential on producing more holistic representations, as our proposed framework generally surpasses state-of-the-art SSL speech/audio models and has superior performance on various datasets compared with other teams in HEAR Challenge. Our code is available at <https://github.com/tony10101105/HEAR-2021-NeurIPS-Challenge—NTU-GURA>.

Keywords: Self-Supervised Learning, Representation Learning, Ensemble Learning

* with equal contribution

1. Introduction

Data representation is crucial to the performance of deep learning algorithms. A good representation catches the underlying patterns of the input data (Bengio et al., 2013) and can serve as effective inputs for neural networks, which usually reach better performance with a higher quality of inputs. The extraction of data can be done either in a hand-crafted way using heuristic methods or with the help of neural networks. For neural-network-based feature extractions, supervised learning, though with a clearer training objective, has limited scalability, while unsupervised learning needs only unlabeled data but usually with inferior results.

In recent years, self-supervised learning (SSL) frameworks for speech representation, such as HuBERT (Hsu et al., 2021) and wav2vec 2.0 (Baevski et al., 2020b), have proved their powerful abilities to extract useful speech audio features and deal with a wide range of downstream tasks, such as Speaker Identification (SID) and Automatic Speaker Verification (ASV). The achievements are mainly acquired by the use of various pre-training techniques with self-defined tasks on a huge amount of unlabeled corpus. For example, during the pre-training stage, HuBERT is required to predict the frame-level cluster assignments of the masked part of the input sequence. By solving self-defined tasks, SSL speech models are capable of capturing meaningful latent representations of given speech audio clips to help with diverse downstream tasks.

Former researches about SSL speech models mainly focus on their representation abilities on speech corpus, while their effectiveness on non-speech tasks still lacks exploration. Also, given that they are all powerful feature extractors and features generated by each of them may be distinct due to different pre-training techniques, studies about speech representation ensemble are not abundant. Therefore, this work proposes a speech representation ensemble framework that integrates several models’ output features with aggregation and concatenation techniques. Our framework is evaluated on the official benchmark of NeurIPS 2021 HEAR Challenge (Turian et al., 2022), which contains 16 speech and non-speech datasets, to investigate both SSL models’ representation abilities on audio data and the feasibility of speech representation ensemble. We found that wav2vec 2.0 and HuBERT are capable of tackling some tasks that are unrelated to their pretraining corpus. For example, HuBERT acquires the second-best result among all submitted single models on Gunshot (Cooper and Shaw, 2020) dataset, which is a sound distance classification task. This indicates HuBERT’s capacity of preserving the information of distance from the sound source. However, we also note that, though wav2vec 2.0 and HuBERT do well on instrument sound classification, their performances on fine-grained music tasks such as pitch classification drop drastically. We also observe the performance gaps between wav2vec 2.0 and HuBERT on various speech and non-speech datasets (the former may have high-quality feature presentations on certain datasets, where the latter cannot extract well, and vice versa). Finally, we demonstrate that, with appropriate representation ensemble strategies, the fusion of different model embeddings, with more holistic information, can achieve higher performances, and our ensemble framework indeed displays competitive results among wav2vec 2.0, HuBERT, and other teams’ approaches.

In summary, the contributions of this work are twofold:

1. Investigating SSL speech models’ effectiveness on diverse audio datasets. We found that current SOTA SSL speech models (wav2vec 2.0 and HuBERT) are capable of some non-speech scenarios, such as instrument sound classification and environmental sound event detection. However, failure on the pitch- and note-related tasks is also observed. In addition, we found the performance gaps between wav2vec 2.0 and HuBERT on several datasets.
2. Proposing a speech representation ensemble framework for integrating SSL speech representation models. Experimental results on HEAR benchmark demonstrate its superiority over state-of-the-art methods and ensemble learning’s potential in the field of speech representation learning.

2. Related Works

2.1. SSL speech and audio Models

The objective of SSL is to extract meaningful representations from the input data without any human annotations. It arranges auxiliary tasks with unlabeled data, driving models to learn the representative features by solving these tasks. Models pre-trained through SSL skills on diverse large corpus can be used as effective feature extractors afterward. As a whole, pre-training strategies for speech/audio models can be roughly categorized into three groups: generative, discriminative, and multi-task learning.

Generative training investigates the distribution of input data. Given a sequence of samples, models are asked to generate new samples to fit into the original distribution. APC (Chung et al., 2019) is pre-trained with a unidirectional RNN in an autoregressive manner to predict the information of future frames. VQ-APC (Chung et al., 2020) extends APC by applying Vector Quantized (VQ) layers to generate better quantized representation. PANNs (Kong et al., 2020), which is an SSL audio model pre-trained on AudioSet (Gemmeke et al., 2017), proposes Wavegram-Logmel-CNN to concatenate both waveform and the log-mel spectrogram as input representations. PaSST (Koutini et al., 2022) presents the patchout method to optimize Transformers with audio spectrograms. Wav2CLIP (Wu et al., 2022) distills representations from CLIP (Radford et al., 2021) to produce general and robust audio representations.

Discriminative training aims to discriminate positive samples from negative samples in the embedding space. CPC (Oord et al., 2018) extracts features that maximally preserve the mutual information of the input sequence and the context latent representations over the time horizon. It makes use of a non-linear CNN-based encoder to embed the input sequence into latent representations and an autoregressive RNN-based network to generate the context latent representations by summarizing the encoder’s previous timesteps. Wav2vec (Schneider et al., 2019) improves CPC by adopting a CNN-based autoregressive network to parallelize the training process and thus enhance time efficiency. vq-wav2vec (Baevski et al., 2020a) further boosts wav2vec by introducing VQ modules, such as Gumbel-Softmax (Jang et al., 2017) and online k-means clustering, to learn discrete speech representations. wav2vec 2.0 refines the framework of vq-wav2vec by merging the two-stage pipeline to perform end-to-end joint training. HuBERT is pre-trained on the task of predicting the

k-means cluster of the masked tokens, currently achieving state-of-the-art performance on many speech tasks.

Multi-task learning solves multiple speech/audio tasks simultaneously to guide models to generate more comprehensive and robust representations. PASE (Pascual et al., 2019) utilizes regressors to predict training objectives, such as waveform, log power spectrum (LPS), mel-frequency cepstral coefficients (MFCC), and prosody features. Discriminators are also introduced into PASE, and strategies to sample positive and negative samples, such as local info max (LIM), global info max (GIM), and sequence predicting coding (SPC) are adopted. PASE+ (Ravanelli et al., 2020) improves PASE by applying online speech distortion modules to add several noises to input speech and using Quasi-RNN (QRNN) (Bradbury et al., 2017) to help the encoder learn the long-term dependencies.

2.2. Ensemble Learning

Ensemble learning is a machine learning technique that aggregates diverse base models to form an ensemble system with better robustness and performance. Generally, ensemble strategies can be categorized into boosting, bagging, and stacking (Ganaie et al., 2022). Boosting strategy integrates weak models, converting them into a new learner with better generalization. Bagging strategy works by training several weak models on datasets where each data point is randomly sampled from the original dataset, aiming to reduce model variance. Stacking strategy takes predictions generated by each model as input, learning how to combine input predictions to achieve better performance. For the feature ensemble of SSL speech representation models, there are some intuitive operations (Arunkumar et al., 2022) such as feature summation, averaging, concatenation, and fusing with attention layers. It has been shown that feature ensemble achieves higher performance in automatic speaker recognition (ASR) (Arunkumar et al., 2022).

3. Methods

We argue that speech models with different SSL techniques produce inhomogeneous feature representations given the same audio input. Hence, even though these cutting-edge models’ performances on various benchmark datasets are close, it could still be beneficial to merge their representations for climbing up to a higher performance. Hence, we propose an ensemble framework that integrates the representations of speech models, which is further shown to be more holistic and contain more information compared with a single SSL speech/audio model.

As shown in Figure 1, the advocated ensemble framework comprises three models and two ensemble techniques. The three models are wav2vec 2.0, HuBERT, and CREPE (Kim et al., 2018b). The two techniques are feature aggregation and feature concatenation, where the former is an intra-model operation that averages the output of a network’s different layers to form a single output feature and the latter is an inter-model operation that combines models’ output features as the final representation of input audio.




Figure 1: The model architecture of the proposed framework. We first take the average of each layer’s feature. Next, linear interpolation is applied to align each constituent model’s averaged output. Finally, representations are concatenated to form the final representation.

3.1. Models for Ensemble

Although the number of models for the bagging method is unlimited, three or five are often chosen in practice for balancing size and performance. Our method in HEAR adopts three models: wav2vec 2.0, HuBERT, and CREPE, which are trained on different datasets, paradigms, and techniques. Wav2vec 2.0 and HuBERT are both SSL speech models, while the former is pre-trained with masked vector quantization plus contrastive discrimination, and the latter is masked vector quantization plus token prediction/classification. Wav2vec 2.0 has Base and Large versions, and HuBERT has Base, Large, and Extra Large versions. The Base versions of both models are pre-trained on LibriSpeech 960 hr (Panayotov et al., 2015), while other versions are trained on Libri-Light 60k hr (Kahn et al., 2020). CREPE is a CNN-based pitch estimation model pre-trained on several pitch-based datasets, such as RWC-Synth (Goto et al., 2002), MDB-STEM-Synth (Salamon et al., 2017) and NSynth, with supervised training manner. Following its official implementation, we use CREPE’s fifth max-pooling layer output as its feature representation of input audio. The adoption of CREPE is based on the observation that HuBERT and wav2vec 2.0 seem unable to handle pitch-related tasks well, which is further discussed in the experiment Section 5.

It should be emphasized that our framework does not limit the number and types of models to be assembled. It is completely plausible to replace wav2vec 2.0, HuBERT, and CREPE with other models, such as vq-wav2vec and TERA, while only the three models are studied during the challenge.

3.2. Feature Aggregation

Each layer of SSL Transformer-based speech models may have attention to a certain aspect of sequential input data. As a result, instead of merely using the last hidden state of the model, we take the average of all the hidden states, i.e., fuse the output of each transformer block plus the initial embedding outputs, to yield more comprehensive feature representa-

tions, as shown in Figure 1. There have been various layer aggregation methods (Yu et al., 2018) studied, while we simply utilize the most straightforward one for simplicity. In other words, though the choice and weight of hidden states are both optional, we sum up all the hidden states fairly. Our aggregation operation can be generally defined as:

$$X = \frac{1}{L} \sum_{i=1}^L X_i, \quad (1)$$

where L is the number of Transformer blocks and X_i is the i -th Transformer block’s output. Notice that the dimensions of different hidden states are all identical in both wav2vec 2.0 and HuBERT, so we can directly take the average, as shown in Equation (1).

It should be noted that feature aggregation is not applied to CREPE since it is a relatively shallow vanilla CNN model with six convolutional layers. We only leverage the output of the fifth max-pooling layer as the feature embedding, following its official implementation*.

3.3. Feature Concatenation

The dimension of a representation model’s output features can be roughly described as $T * C$, where T is the timestamp dimension and C is the feature dimension of each timestamp. Since HuBERT, wav2vec 2.0, and CREPE have different output dimensions (the output dimension of HuBERT and wav2vec 2.0 depends on their versions, such as wav2vec 2.0 Base or wav2vec 2.0 Large), up/downsampling is required for feature concatenation. For simplicity, we linearly interpolate the timestamp dimension T and feature dimension C of each constituent model’s representation into those of wav2vec 2.0. Next, features are concatenated along the feature dimension so that the dimension of the final output feature is $T_w * (3C_w)$, where the subscript w denotes wav2vec 2.0. This process preserves the time information of the final output feature, which is vital to certain downstream tasks such as sound event detection.

4. Experimental Setup

Experiments are carried out during the NeurIPS 2021 HEAR Challenge, where 16 iconic or novel datasets are used for the holistic evaluation of audio representations. HEAR officials conduct the evaluation process, and the full results are displayed on its official benchmark website.†

4.1. HEAR Challenge and Evaluation Pipeline

HEAR Challenge is one of the challenges in NeurIPS 2021 which aims at developing audio representation frameworks with high generalizability. HEAR has established a comprehensive benchmark that currently includes 19 downstream tasks, including but not limited to speech, environmental sound, and music domains. Our experiments follow HEAR’s evaluation pipeline, which adopts the evaluation principles of representation quality proposed

*<https://github.com/marl/crepe>

†<https://hearbenchmark.com>

Table 1: Categories of 16 datasets presented in HEAR Challenge’s evaluation process.

Category	Datasets
Speech	CREMA-D (Cao et al., 2014) LibriCount (Stöter et al., 2018) Speech Commands (Warden, 2018) VoxLingua107 (Valk and Alumäe, 2021) Vocal Imitation Set (Kim et al., 2018a)
Instrument	Nsynth (Engel et al., 2017) MAESTRO (Hawthorne et al., 2019) Beijing Opera Percussion (Tian et al., 2014) Mridingham Stroke and Tonic (Anantapadmanabhan et al., 2013)
Environmental Sounds	ESC-50 (Piczak, 2015) Beehive States (Nolasco et al., 2019) DCASE 2016 Task 2 (Mesaros et al., 2018) Gunshot Triangulation (Cooper and Shaw, 2020)
General	FSD50k (Fonseca et al., 2022) GTZAN Music/Speech (Tzanetakis, 1999) GTZAN Genre Collection (Tzanetakis and Cook, 2002)

by (Goyal et al., 2019). In brief, HEAR evaluates the effectiveness of submitted upstream models by investigating the performances of downstream shallow models, trained with embeddings generated by upstream models, on each task. Since HEAR’s philosophy is to foster the development of audio representation methods with strong generalizability, models are anticipated to have decent performances on all datasets, rather than excellent results on several datasets but poor on other tasks.

Notably, HEAR requires challenges to provide two representations: timestamp embedding and scene embedding. The former is the representation that has one dimension as the time dimension, and the latter is a single embedding that represents the whole audio. A baseline method to generate scene embedding is to take the average of timestamp embedding along the time dimension. This work adopts this approach during the challenge.

4.2. Datasets

HEAR benchmark currently contains 19 tasks with 16 datasets. Some tasks utilize the same datasets. Datasets adopted in HEAR include Speech Commands (Warden, 2018), NSynth (Engel et al., 2017), DCASE 2016 Task 2 (Mesaros et al., 2018), Beehive States (Nolasco et al., 2019), Beijing Opera Percussion Instrument (Tian et al., 2014), CREMA-D (Cao et al., 2014), ESC-50 (Piczak, 2015), FSD50k (Fonseca et al., 2022), Gunshots recorded in an open field using iPod Touch devices (Gunshot Triangulation) (Cooper and Shaw, 2020), GTZAN Genre Collection (Tzanetakis and Cook, 2002), GTZAN Music/Speech (Tzanetakis, 1999), LibriCount (Stöter et al., 2018), MAESTRO (Hawthorne et al., 2019), Mridangam Stroke (Anantapadmanabhan et al., 2013), Vocal Imitation Set (Kim et al., 2018a) and VoxLingua107 (Valk and Alumäe, 2021).

Categories of these 16 datasets are presented in Table 1. Speech Command is a speech command classification dataset that contains 105829 utterances of 35 words. NSynth is a pitch classification dataset from 1006 instruments. DCASE 2016 Task 2 is an office sound

event detection dataset with 11 event categories, such as clearing throat and coughing. It should be noticed that HEAR uses a different way to split this dataset, so the evaluation result is not comparable to other work. Beehive States is a binary beehive sound classification dataset with 930 sound clips. Beijing Opera Percussion Instrument is a percussion instrument classification dataset with 6 percussion instruments, such as Ban and Gong, which are organized into four classes, such as Bangu and Daluo. CREMA-D is a 6-class audiovisual emotion recognition dataset. HEAR only adopts its audio recordings. ESC-50 is a 50-class environmental sounds classification dataset. FSD50K is a 200-class sound event classification dataset with 51197 Freesound (Fonseca et al., 2017) clips drawn from AudioSet. Gunshot Triangulation is a sound source distance classification dataset. GTZAN Genre Collection is a 10-class music genre classification dataset. GTZAN Music/Speech is a speech-or-music binary classification dataset. LibriCount is a speaker number estimation dataset. MAESTRO is a note onset detection dataset. The Mridangam Stroke dataset contains 6977 audio examples of 10 different strokes played on Mridangams with 6 tonics. HEAR evaluation divides the dataset into two sub-tasks: stroke classification and tonic classification. The Vocal Imitation Set is a sound classification dataset. Models should predict which reference sound the audio is imitating. VoxLingua107 is a spoken-language classification dataset. HEAR adopts the 10-language subset.

4.3. Our Method and Baselines

Our ensemble framework adopts pre-trained wav2vec 2.0, HuBERT, and CREPE without any fine-tuning. As a result, the utilized wav2vec 2.0 and HuBERT are both pure speech models pre-trained on LibriSpeech or Libri-Light, while CREPE is a pure music-based model pre-trained on several pitch-related datasets.

Due to feature concatenation, our proposed ensemble method generates longer scene embeddings than the single SSL model. This affects the architecture of interfaced downstream model since its input layer should match the size of input scene embeddings. Hence, it is necessary to ensure our framework does not gain benefits from the architectures of downstream models. To investigate this issue, we design a variant of the baseline scene embeddings generation approach that first splits the timestamp embeddings along the time dimension into k groups, takes the mean along the time dimension, and concatenates each sub-scene embeddings as final scene embeddings. Thus, the length of scene embeddings will be $(k * C_w)$ rather than C_w . With this pipeline, the scene embeddings are much longer, while the representation information of the three models will still be preserved. The group number k is set to 5.

Baselines include vanilla wav2vec 2.0, HuBERT, CREPE, and the other five teams' methods. For vanilla wav2vec 2.0 and HuBERT, only the last hidden state is used as the input embedding. The five teams' models are officially considered to have the most competitive performances out of all participants. Five models include Wav2CLIP (Wu et al., 2022), PaSST (Koutini et al., 2022), PANNs (Kong et al., 2020), EfficientNet-B2 (Tan and Le, 2019) and BYOL-S (Scheidwasser-Clow et al., 2022). The first three are SSL speech and audio models, while EfficientNet is a series of CNN-based models whose architectures are determined by neural architecture search (NAS), denoted as EfficientNet-B0 to EfficientNet-B7. BYOL-S is proposed in SERAB (Scheidwasser-Clow et al., 2022) benchmark as a speech

version of BYOL (Grill et al., 2020). BYOL is an image representation learning framework that comprises two networks with the same architectures, referred to as the online network and target network. By taking an image as the online network’s input and the augmented view as the target network’s input, BYOL aims to minimize the distance between the output embeddings of these two networks. For the type of training corpus, Wav2CLIP is pre-trained on VGG-Sound (Chen et al., 2020). PaSST, BYOL-S, and PANNs are pre-trained on AudioSet. EfficientNet-B2 is pre-trained on ImageNet (Deng et al., 2009). Notably, each team’s model implementation may be slightly different from the original paper of their methods.

5. Results

5.1. Quantitative Results

The results of 18 tasks containing 15 out of 16 datasets are displayed in Table 2. We exclude the Beehive States dataset since it has a large input size which causes most teams’ models to run out of memory under the competition hardware setting and get no result. The highest and second-highest scores in each dataset are marked in bold and underlined respectively.

It is demonstrated that `fusion_cat_xwc` possesses competitive performances on a large portion of datasets presented in Table 1. Specifically, our method achieves top performances on Beijing Opera Percussion, CREMA-D, MAESTRO 5h, Mridingham Stroke, Speech Commands 5h, Speech Commands full, Vocal Imitations, and VoxLingua107, along with second places on DCASE 2016 Task 2, Gunshot Triangulation, LibriCount, Mridingham Tonic, NSynth 5h, and NSynth 50h. The datasets our method does not reach the top two results are ESC-50, FSD50k, GTZAN Genre, and GTZAN Music/Speech. It can also be observed that our ensemble method outperforms the state-of-the-art HuBERT xLarge and wav2vec 2.0 Large alone on all datasets except GTZAN Music/Speech and LibriCount, indicating that the ensemble framework can indeed stably increase a single model’s speech/audio representation ability by incorporating multiple representation models. Furthermore, the `fusion_cat_xwc_g` with much longer scene embeddings does not surpass the performance of the original `fusion_cat_xwc`. Specifically, it gains lower performances on 9 tasks, slightly higher results on 6 tasks, and 3 ties on tasks that only require timestamp embeddings. This shows that longer scene embeddings do not guarantee a stronger downstream model and thus higher downstream performances.

In terms of other submitted methods, PaSST base2levelmel owns first place on 6 tasks and second place on 1 dataset, which is the second-best model in general. CREPE is exceptionally strong at pitch- and note-related music tasks, such as NSynth 5h, NSynth 50h, and MAESTRO 5h, which most current speech/audio representation models fail to cope with. However, the authors note that CREPE’s poor performances on other datasets can also be found. Wav2vec 2.0 and HuBERT, though only pretrained on speech corpus, are effective on various speech and non-speech tasks. In particular, for non-speech tasks, except our proposed ensemble framework, HuBERT achieves the second-best results on Gunshot Triangulation, Mridingham Tonic, and NSynth 5h. Nevertheless, their performances on the pitch- and note-related music tasks are not satisfying. Both wav2vec 2.0 and HuBERT fail to deal with MAESTRO 5h, and their performances on NSynth also substantially fall behind CREPE. Especially, HuBERT only gets 18.4 accuracy on NSynth 5h and 42.9 on

Table 2: Performance of our framework and other baselines on 18 tasks. Some datasets are only partially evaluated, as described in Section 4. Our method is denoted as fusion_cat_xwc, which means HuBERT xLarge plus wav2vec 2.0 Large plus CREPE with intra-model feature aggregation and inter-model feature concatenation. The group-based scene embedding generation technique with our framework is fusion_cat_xwc_g. Note that fusion_cat_xwc generally outperforms all of its constituent models and has competitive performances among all teams’ approaches.

Task and measure	Task and measure								
	Beijing Opera (Acc. ↑)	CREMA-D (Acc. ↑)	DCASE 2016 (Onset FMS ↑)	ESC-50 (Acc. ↑)	FSD50k (mAP ↑)	GTZAN Genre (Acc. ↑)	GTZAN Music/Speech (Acc. ↑)	Gunshot (Acc. ↑)	Libricount (Acc. ↑)
PaSST base2levelmel (Koutini et al., 2022)	96.6	61.0	92.5	94.7	64.1	88.3	<u>97.7</u>	94.0	66.0
PANNS (Kong et al., 2020)	91.1	55.5	00.0	90.9	-	86.0	99.2	79.8	65.2
Wav2CLIP (Wu et al., 2022)	93.6	51.2	00.0	75.9	36.2	74.8	94.6	92.9	52.8
SERAB BYOL-S (Scheidwasser-Clow et al., 2022)	95.3	65.7	64.2	80.5	50.9	83.7	93.8	85.7	78.5
EfficientNet-B2 (Tan and Le, 2019)	95.3	57.5	79.0	<u>93.5</u>	<u>60.7</u>	<u>87.8</u>	96.8	87.8	65.1
CREPE (Kim et al., 2018b)	92.8	38.3	50.4	30.0	15.9	64.5	92.9	86.3	49.9
wav2vec 2.0 (Baevski et al., 2020b)	90.7	65.6	66.3	56.1	34.2	78.0	94.6	84.8	69.2
HuBERT xLarge (Hsu et al., 2021)	94.5	69.0	58.4	60.3	31.4	73.5	91.3	93.2	64.6
fusion_cat_xwc_g	<u>96.2</u>	<u>74.3</u>	<u>82.6</u>	65.3	37.4	76.0	94.4	90.5	65.9
fusion_cat_xwc	96.6	74.7	<u>82.6</u>	73.4	42.0	80.5	92.8	<u>93.5</u>	<u>69.7</u>

Task and measure	Task and measure								
	Maestro 5h (Onset FMS ↑)	Mridingham Stroke (Acc. ↑)	Mridingham Tonic (Acc. ↑)	NSynth Pitch 50h (Pitch Acc. ↑)	NSynth Pitch 5h (Pitch Acc. ↑)	Speech commands 5h (Acc. ↑)	Speech commands full (Acc. ↑)	Vocal Imitation (mAP ↑)	VoxLingual107 top 10 (Acc. ↑)
PaSST base2levelmel (Koutini et al., 2022)	-	96.5	81.9	54.1	25.6	68.1	63.9	18.2	25.9
PANNS (Kong et al., 2020)	00.0	93.9	82.4	30.1	14.8	56.0	61.8	12.7	24.4
Wav2CLIP (Wu et al., 2022)	00.0	94.7	82.9	43.9	23.0	31.6	34.7	08.3	19.2
SERAB BYOL-S (Scheidwasser-Clow et al., 2022)	00.8	<u>97.3</u>	92.8	71.2	39.6	91.4	94.8	16.0	45.8
EfficientNet-B2 (Tan and Le, 2019)	00.0	94.9	84.3	39.1	16.8	57.3	67.6	13.8	25.5
CREPE (Kim et al., 2018b)	40.1	89.8	82.4	90.0	87.0	18.0	21.1	05.1	14.2
wav2vec 2.0 (Baevski et al., 2020b)	03.3	94.3	82.8	65.3	40.2	83.8	87.9	08.0	49.3
HuBERT xLarge (Hsu et al., 2021)	00.7	95.3	85.0	42.9	18.4	<u>95.3</u>	<u>95.4</u>	15.4	<u>63.7</u>
fusion_cat_xwc_g	44.1	97.5	<u>92.4</u>	<u>89.1</u>	<u>85.4</u>	95.1	96.8	21.5	62.9
fusion_cat_xwc	44.1	97.2	92.3	88.5	84.6	96.1	96.8	<u>19.7</u>	72.0

NSynth 50h, which are the third from last among all methods. On the contrary, wav2vec 2.0 has over 20% higher results than HuBERT on this task. This shows that representation models’ feature extraction abilities on non-speech data are also affected by the design of their pretraining schema beside the training corpus.

It is also noteworthy that, though AudioSet and VGG-Sound are two large-scale datasets containing both speech and non-speech audio clips, models pre-trained on them, such as PaSST, PANNs, and Wav2CLIP, do not remarkably surpass other models that are only pre-trained on speech audio, such as SERAB BYOL-S, HuBERT xLarge and wav2vec 2.0 Large, or even images, such as EfficientNet-B2. In addition, models pre-trained on the same corpus can still have great performance gaps on testing datasets. For instance, PaSST base2levelmel acquires 54.1 Pitch accuracy (Acc) on NSynth 50h, while PANNs only get 30.1; wav2vec 2.0 Large has 65.3 Pitch accuracy on NSynth 50h, yet HuBERT xLarge only gets 42.9. This leads to a conclusion that, besides pre-training corpus, models’ attention to input audio can also be influenced by other factors, such as model architectures and pre-training objectives. Even though existing SSL speech models like wav2vec 2.0 and HuBERT normally have similar Transformer-based structures, their abilities to extract speech and non-speech audio can still vary. This observation turns out to be the basis of the use of ensemble learning, and the promising performances of our proposed ensemble framework in turn validate this observation.

5.2. Ablation Study

Ablation studies are conducted to verify the contributions of operations and components in our ensemble framework. Results are shown in Tables 3, 4, 5 and 6.

We first look into the use of intra-model feature aggregation. As displayed in Table 3, the direct fusion of all hidden states helps models generate better representations. For instance, for DCASE 2016 Task 2, the event onset f-measure (FMS) of wav2vec 2.0 Large increases from 66.3 to 79.8 with the fusion strategy. HuBERT xLarge follows the same trend but with more significant improvements, with event onset FMS rocketing to 82.6. However, feature aggregation does not always produce better results. For example, wav2vec 2.0 Large’s classification accuracy on LibriCount drops from 69.2 to 65.3 after applying the feature fusion. Nevertheless, this technique still generally has a positive effect on models’ representation ability. As for feature concatenation versus feature averaging, Table 3 shows that the cat xw has better performances on eleven tasks and avg xw wins seven, with one tie. With the abundant results shown on the 18 tasks, we conclude that, by and large, feature concatenation brings more benefits than feature averaging, and so we adopt the former to our framework. Table 5 illustrates the effects of adding CREPE. With CREPE, HuBERT xLarge’s performances on DCASE 2016 Task 2, GTZAN Music/Speech, MAESTRO 5h, NSynth 50h, and NSynth 5h rise, especially significantly in NSynth 50h, NSynth 5h, and MAESTRO 5h. The four tasks are all music-related. Besides, DCASE 2016 Task 2 and MAESTRO 5h are event onset detection datasets. It can be derived that, SSL speech models may still have blind spots in music-related content, and their abilities on event detection datasets are also not as strong as in other kinds of tasks such as classification/identification. Though incorporating CREPE solves the issue, it leads to drops in other tasks by a considerable margin. For instance, its performance on Speech Commands full decreases from 95.4

Table 3: Performance of HuBERT xLarge and wav2vec 2.0 Large with and without feature fusion. Direct feature fusion over all hidden states increases SSL speech models’ performances on most tasks, including speech and non-speech ones.

	Task and measure								
	Beijing Opera (Acc. ↑)	CREMA-D (Acc. ↑)	DCASE 2016 (Onset FMS ↑)	ESC-50 (Acc. ↑)	FSD50k (mAP ↑)	GTZAN Genre (Acc. ↑)	GTZAN Music/Speech (Acc. ↑)	Gunshot (Acc. ↑)	Libricount (Acc. ↑)
wav2vec 2.0 Large	90.7	65.6	66.3	56.1	34.2	78.0	94.6	84.8	69.2
Fusion wav2vec 2.0 Large	94.5	69.2	79.8	69.5	40.3	79.3	95.3	96.7	65.3
HuBERT xLarge	94.5	69.0	58.4	60.3	31.4	73.5	91.3	93.2	64.6
Fusion HuBERT xLarge	94.9	75.2	82.6	74.3	41.3	79.6	93.6	92.9	68.3
	Task and measure								
	Maestro 5h (Onset FMS ↑)	Mridingham Stroke (Acc. ↑)	Mridingham Tonic (Acc. ↑)	NSynth Pitch 50h (Pitch Acc. ↑)	NSynth Pitch 5h (Pitch Acc. ↑)	Speech commands 5h (Acc. ↑)	Speech commands full (Acc. ↑)	Vocal Imitation (mAP ↑)	VoxLingua107 top 10 (Acc. ↑)
wav2vec 2.0 Large	03.3	94.3	82.8	65.3	40.2	83.8	87.9	08.0	49.3
Fusion wav2vec 2.0 Large	11.1	96.2	83.8	60.6	33.0	95.7	96.9	17.4	70.6
HuBERT xLarge	00.7	95.3	85.0	42.9	18.4	95.3	95.4	15.4	63.7
Fusion HuBERT xLarge	16.6	97.4	90.9	68.8	38.2	94.7	95.7	18.5	71.4

Table 4: Performance of feature averaging versus feature concatenation of wav2vec 2.0 Large plus HuBERT Large. Cat xw and avg xw denote the concatenation and the average of HuBERT Large and wav2vec 2.0 Large’s representations, respectively.

Task and measure	Beijing Opera (Acc. \uparrow)	CREMA-D (Acc. \uparrow)	DCASE 2016 (Onset FMS \uparrow)	ESC-50 (Acc. \uparrow)	FSD50k (mAP \uparrow)	GTZAN Genre (Acc. \uparrow)	GTZAN Music/Speech (Acc. \uparrow)	Gunshot (Acc. \uparrow)	Libricount (Acc. \uparrow)
avg xw	94.1	69.9	41.7	58.7	31.8	74.6	92.8	81.0	62.3
cat xw	93.6	69.8	45.2	58.6	32.3	73.4	93.6	86.9	62.7
Task and measure	Maestro 5h (Onset FMS \uparrow)	Mridingham Stroke (Acc. \uparrow)	Mridingham Tonic (Acc. \uparrow)	NSynth Pitch 50h (Pitch Acc. \uparrow)	NSynth Pitch 5h (Pitch Acc. \uparrow)	Speech commands 5h (Acc. \uparrow)	Speech commands full (Acc. \uparrow)	Vocal Imitation (mAP \uparrow)	VoxLingua107 top 10 (Acc. \uparrow)
avg xw	00.4	94.6	79.9	41.5	19.8	90.7	95.2	16.5	69.0
cat xw	00.3	95.1	80.2	42.8	19.8	93.6	92.9	16.6	70.6

Table 5: Performance of HuBERT xLarge with and without CREPE. Avg xc means HuBERT Large plus CREPE with simple feature averaging. Notice that incorporating CREPE can largely improve performances on music-related tasks, such as NSynth and MAESTRO. This demonstrates HuBERT’s failure in music tasks as well as the effectiveness of the representation ensemble.

Task and measure	Beijing Opera (Acc. ↑)	CREMA-D (Acc. ↑)	DCASE 2016 (Onset FMS ↑)	ESC-50 (Acc. ↑)	FSD50k (mAP ↑)	GTZAN Genre (Acc. ↑)	GTZAN Music/Speech (Acc. ↑)	Gunshot (Acc. ↑)	Libricount (Acc. ↑)
HuBERT xLarge	94.5	69.0	58.4	60.3	31.4	73.5	91.3	93.2	64.6
avg xc	93.2	54.0	61.0	43.7	25.3	69.8	94.6	84.5	62.2
Task and measure	Maestro 5h (Onset FMS ↑)	Mirdingham Stroke (Acc. ↑)	Mirdingham Tonic (Acc. ↑)	NSynth Pitch 50h (Pitch Acc. ↑)	NSynth Pitch 5h (Pitch Acc. ↑)	Speech commands 5h (Acc. ↑)	Speech commands full (Acc. ↑)	Vocal Imitation (mAP ↑)	VoxLingua107 top 10 (Acc. ↑)
HuBERT xLarge	00.7	95.3	85.0	42.9	18.4	95.3	95.4	15.4	63.7
avg xc	46.2	91.7	83.1	89.7	87.8	73.7	82.3	07.9	29.3

Table 6: Performance of wav2vec 2.0 Large plus CREPE with feature concatenation with and without HuBERT xLarge. Cat wc denotes the ensemble of wav2vec 2.0 Large plus CREPE with feature concatenation, while cat xwc means the same schema but with the join of HuBERT xLarge. Note that cat xwc generally outperforms cat wc, which means HuBERT xLarge provides extra acoustic information that makes the concatenated feature more holistic.

Task and measure	Beijing Opera (Acc. \uparrow)	CREMA-D (Acc. \uparrow)	DCASE 2016 (Onset FMS \uparrow)	ESC-50 (Acc. \uparrow)	FSD50k (mAP \uparrow)	GTZAN Genre (Acc. \uparrow)	GTZAN Music/Speech (Acc. \uparrow)	Gunshot (Acc. \uparrow)	Libricount (Acc. \uparrow)
cat wc	92.0	46.0	58.5	34.3	23.4	68.1	93.8	83.3	56.9
cat xwc	96.6	74.2	82.6	73.4	42.0	80.5	92.8	93.5	69.7
Task and measure	Maestro 5h (Onset FMS \uparrow)	Mridingham Stroke (Acc. \uparrow)	Mridingham Tonic (Acc. \uparrow)	NSynth Pitch 50h (Pitch Acc. \uparrow)	NSynth Pitch 5h (Pitch Acc. \uparrow)	Speech commands 5h (Acc. \uparrow)	Speech commands full (Acc. \uparrow)	Vocal Imitation (mAP \uparrow)	VoxLingua107 top 10 (Acc. \uparrow)
cat wc	46.3	89.8	82.3	89.9	86.8	88.5	91.9	07.6	31.0
cat xwc	44.1	97.2	92.3	88.5	84.6	96.1	96.8	19.7	72.0

accuracy to 82.3 after combining CREPE. The reason is that, given that embeddings on Speech Command full produced by CREPE contain very little content information, it will further jeopardize the information in embeddings generated by HuBERT Large if we take the average of these two vectors. The solution is to use feature concatenation rather than averaging, and our concatenation-based final framework indeed achieves 96.8 accuracy on Speech Command full even with CREPE being inside. Finally, the contribution of HuBERT xLarge among all three models is examined in Table 6. It is shown that HuBERT xLarge helps cat xwc attain higher results on nearly all tasks. For example, the event onset FMS rises from 58.5 to 82.6 in DCASE 2016 Task 2, and accuracy from 34.3 to 73.4 in ESC-50.

6. Conclusion

This work proposes a representation ensemble framework that integrates models’ audio representations and leverages the framework to investigate SSL speech models’ effectiveness on speech and non-speech datasets. Experimental results show that SSL speech models’ are capable of extracting meaningful representations of various non-speech corpus, such as instrument classification, while failure on fine-grained music-based tasks, such as pitch- and note-related datasets is also observed. In addition, different SSL speech models might have insights into different aspects of audio features. Specifically, wav2vec 2.0 and HuBERT are shown to have considerable performance gaps on some datasets, such as NSynth Pitch and VoxLingua107. We also investigate the contribution of each component, such as feature concatenation, of the ensemble framework, and the overall performance of our method generally surpasses its state-of-the-art constituent models and other teams’ methods in HEAR Challenge.

References

- Akshay Anantapadmanabhan, Ashwin Bellur, and Hema A Murthy. Modal analysis and transcription of strokes of the Mridangam using non-negative matrix factorization. In *Proceedings of the 38th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 181–185, 2013.
- A Arunkumar, Vrunda Nileshkumar Sukhadia, and Srinivasan Umesh. Investigation of Ensemble features of Self-Supervised Pretrained Models for Automatic Speech Recognition. In *Proceedings of the 23th Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pages 5145–5149, 2022. doi: 10.21437/Interspeech.2022-11376.
- Alexei Baevski, Steffen Schneider, and Michael Auli. vq-wav2vec: Self-supervised learning of discrete speech representations. In *Proceedings of the 8th International Conference on Learning Representations (ICLR)*, 2020a.
- Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, 2020b.

- Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- James Bradbury, Stephen Merity, Caiming Xiong, and Richard Socher. Quasi-recurrent neural networks. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, 2017.
- Houwei Cao, David G Cooper, Michael K Keutmann, Ruben C Gur, Ani Nenkova, and Ragini Verma. CREMA-D: Crowd-sourced Emotional Multimodal Actors Dataset. *IEEE Transactions on Affective Computing*, 5(4):377–390, 2014.
- Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. Vggsound: A large-scale audio-visual dataset. In *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2020.
- Yu-An Chung, Wei-Ning Hsu, Hao Tang, and James Glass. An Unsupervised Autoregressive Model for Speech Representation Learning. In *Proceedings of the 20th Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pages 146–150, 2019. doi: 10.21437/Interspeech.2019-1473.
- Yu-An Chung, Hao Tang, and James Glass. Vector-Quantized Autoregressive Predictive Coding. In *Proceedings of the 21th Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pages 3760–3764, 2020. doi: 10.21437/Interspeech.2020-1228.
- Seth Cooper and Steven Shaw. Gunshots recorded in an open field using iPod Touch devices, 2020. URL <https://doi.org/10.5061/dryad.wm37pvmkc>.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009.
- Jesse H. Engel, Cinjon Resnick, Adam Roberts, Sander Dieleman, Mohammad Norouzi, Douglas Eck, and Karen Simonyan. Neural Audio Synthesis of Musical Notes with WaveNet Autoencoders. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning (ICML)*, volume 70 of *Proceedings of Machine Learning Research*, pages 1068–1077, 2017.
- Eduardo Fonseca, Jordi Pons, Xavier Favory, Frederic Font, Dmitry Bogdanov, Andrés Ferraro, Sergio Oramas, Alastair Porter, and Xavier Serra. Freesound datasets: A platform for the creation of open audio datasets. In *International Society for Music Information Retrieval Conference (ISMIR)*, 2017.
- Eduardo Fonseca, Xavier Favory, Jordi Pons, Frederic Font, and Xavier Serra. FSD50K: An Open Dataset of Human-Labeled Sound Events. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30:829–852, 2022. doi: 10.1109/TASLP.2021.3133208.

- M.A. Ganaie, Minghui Hu, A.K. Malik, M. Tanveer, and P.N. Suganthan. Ensemble deep learning: A review. *Engineering Applications of Artificial Intelligence*, 115: 105151, 2022. doi: <https://doi.org/10.1016/j.engappai.2022.105151>. URL <https://www.sciencedirect.com/science/article/pii/S095219762200269X>.
- Jort F Gemmeke, Daniel P W Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. Audio Set: An ontology and human-labeled dataset for audio events. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 776–780, 2017.
- Masataka Goto, Hiroki Hashiguchi, Takuichi Nishimura, and Ryu ichi Oka. Rwc music database: Popular, classical and jazz music databases. In *International Society for Music Information Retrieval Conference (ISMIR)*, 2002.
- Priya Goyal, Dhruv Mahajan, Abhinav Gupta, and Ishan Misra. Scaling and Benchmarking Self-Supervised Visual Representation Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6390–6399, 2019. URL <https://github.com/facebookresearch/>.
- Jean-Bastien Grill, Florian Strub, Florent Alché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, koray kavukcuoglu, Remi Munos, and Michal Valko. Bootstrap your own latent - a new approach to self-supervised learning. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 21271–21284. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/f3ada80d5c4ee70142b17b8192b2958e-Paper.pdf>.
- Curtis Hawthorne, Andriy Stasyuk, Adam Roberts, Ian Simon, Cheng-Zhi Anna Huang, Sander Dieleman, Erich Elsen, Jesse Engel, and Douglas Eck. Enabling Factorized Piano Music Modeling and Generation with the MAESTRO Dataset. In *Proceedings of the 7th International Conference on Learning Representations (ICLR)*, 2019.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. HuBERT: Self-Supervised Speech Representation Learning by Masked Prediction of Hidden Units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460, 2021.
- Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, 2017.
- Jacob Kahn, Morgane Rivière, Weiyi Zheng, Evgeny Kharitonov, Qiantong Xu, Pierre-Emmanuel Mazaré, Julien Karadayi, Vitaliy Liptchinsky, Ronan Collobert, Christian Fuegen, et al. Libri-light: A benchmark for asr with limited or no supervision. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7669–7673, 2020.

- Bongjun Kim, Madhav Ghei, Bryan Pardo, and Zhiyao Duan. Vocal imitation set: a dataset of vocally imitated sound events using the AudioSet ontology. In Mark D. Plumbley, Christian Kroos, Juan P. Bello, Gaël Richard, Daniel P. W. Ellis, and Annamaria Mesaros, editors, *Proceedings of the Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, volume 1, pages 148–152, 2018a.
- Jong Wook Kim, Justin Salamon, Peter Li, and Juan Pablo Bello. Crepe: A Convolutional Representation for Pitch Estimation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 161–165, 2018b.
- Qiuqiang Kong, Yin Cao, Turab Iqbal, Yuxuan Wang, Wenwu Wang, and Mark D Plumbley. PANNs: Large-Scale Pretrained Audio Neural Networks for Audio Pattern Recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:2880–2894, 2020.
- Khaled Koutini, Jan Schlüter, Hamid Eghbal-zadeh, and Gerhard Widmer. Efficient Training of Audio Transformers with Patchout. In *Proceedings of the 23th Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pages 2753–2757, 2022. doi: 10.21437/Interspeech.2022-227.
- Annamaria Mesaros, Toni Heittola, Emmanouil Benetos, Peter Foster, Mathieu Lagrange, Tuomas Virtanen, and Mark D Plumbley. Detection and Classification of Acoustic Scenes and Events: Outcome of the DCASE 2016 Challenge. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 26(2):379–393, 2018.
- Inês Nolasco, Alessandro Terenzi, Stefania Cecchi, Simone Orcioni, Helen L Bear, and Emmanouil Benetos. Audio-based Identification of Beehive States. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8256–8260, 2019.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: An ASR corpus based on public domain audio books. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015.
- Santiago Pascual, Mirco Ravanelli, Joan Serra, Antonio Bonafonte, and Yoshua Bengio. Learning Problem-Agnostic Speech Representations from Multiple Self-Supervised Tasks. In Gernot Kubin and Zdravko Kacic, editors, *Proceedings of the 20th Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pages 161–165, 2019.
- Karol J Piczak. ESC: Dataset for Environmental Sound Classification. In Xiaofang Zhou, Alan F. Smeaton, Qi Tian, Dick C. A. Bulterman, Heng Tao Shen, Ketan Mayer-Patel, and Shuicheng Yan, editors, *Proceedings of the 23rd Annual ACM Conference on Multimedia Conference (SIGMM)*, pages 1015–1018, 2015.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and

- Ilya Sutskever. Learning transferable visual models from natural language supervision. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning (ICML)*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR, 2021.
- Mirco Ravanelli, Jianyuan Zhong, Santiago Pascual, Pawel Swietojanski, Joao Monteiro, Jan Trmal, and Yoshua Bengio. Multi-task self-supervised learning for robust speech recognition. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6989–6993. IEEE, 2020.
- Justin Salamon, Rachel M. Bittner, Jordi Bonada, Juan J. Bosch, Emilia Gómez, and Juan Pablo Bello. An analysis/synthesis framework for automatic f0 annotation of multi-track datasets. In *International Society for Music Information Retrieval Conference (ISMIR)*, 2017.
- Neil Scheidwasser-Clow, Mikolaj Kegler, Pierre Beckmann, and Milos Cernak. Serab: A multi-lingual benchmark for speech emotion recognition. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7697–7701, 2022. doi: 10.1109/ICASSP43922.2022.9747348.
- Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli. wav2vec: Unsupervised Pre-Training for Speech Recognition. In Gernot Kubin and Zdravko Kacic, editors, *Proceedings of the 20th Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pages 3465–3469, 2019.
- Fabian-Robert Stöter, Soumitro Chakrabarty, Emanuël Habets, and Bernd Edler. Libri-Count, a dataset for speaker count estimation, 2018. URL <https://doi.org/10.5281/zenodo.1216072>.
- Mingxing Tan and Quoc V Le. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97, pages 6105–6114, 2019.
- Mi Tian, Ajay Srinivasamurthy, Mark Sandler, and Xavier Serra. A study of instrument-wise onset detection in Beijing Opera percussion ensembles. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2159–2163, 2014. doi: 10.1109/ICASSP.2014.6853981. URL <http://dx.doi.org/10.1109/ICASSP.2014.6853981>.
- Joseph Turian, Jordie Shier, Humair Raj Khan, Bhiksha Raj, Björn W. Schuller, Christian J. Steinmetz, Colin Malloy, George Tzanetakis, Gissel Velarde, Kirk McNally, Max Henry, Nicolas Pinto, Camille Noufi, Christian Clough, Dorien Herremans, Eduardo Fonseca, Jesse Engel, Justin Salamon, Philippe Esling, Pranay Manocha, Shinji Watanabe, Zeyu Jin, and Yonatan Bisk. HEAR: Holistic Evaluation of Audio Representations. In Douwe Kiela, Marco Ciccone, and Barbara Caputo, editors, *Proceedings of the NeurIPS 2021 Competitions and Demonstrations Track*, volume 176 of *Proceedings of Machine Learning Research*, pages 125–145. PMLR, 06–14 Dec 2022. URL <https://proceedings.mlr.press/v176/turian22a.html>.

- George Tzanetakis. GTZAN Music/Speech Collection, 1999. URL <http://marsyas.info/index.html>.
- George Tzanetakis and Perry Cook. Musical genre classification of audio signals. *IEEE Transactions on Speech and Audio Processing*, 10(5):293–302, 2002.
- Jörgen Valk and Tanel Alumäe. VOXLINGUA107: A Dataset for Spoken Language Recognition. In *IEEE Spoken Language Technology Workshop (SLT)*, pages 652–658, 2021.
- Pete Warden. Speech Commands: A Dataset for Limited-Vocabulary Speech Recognition. *CoRR arXiv*, 2018. URL <http://arxiv.org/abs/1804.03209>.
- Ho-Hsiang Wu, Prem Seetharaman, Kundan Kumar, and Juan Pablo Bello. Wav2clip: Learning robust audio representations from clip. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4563–4567, 2022. doi: 10.1109/ICASSP43922.2022.9747669.
- Fisher Yu, Dequan Wang, Evan Shelhamer, and Trevor Darrell. Deep layer aggregation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2403–2412, 2018.