



Masked Modeling Duo for Speech: Specializing General-Purpose Audio Representation to Speech using Denoising Distillation

Daisuke Niizumi, Daiki Takeuchi, Yasunori Ohishi, Noboru Harada, and Kunio Kashino

NTT Corporation, Japan

daisuke.niizumi@ntt.com

Abstract

Self-supervised learning general-purpose audio representations have demonstrated high performance in a variety of tasks. Although they can be optimized for application by fine-tuning, even higher performance can be expected if they can be specialized to pre-train for an application. This paper explores the challenges and solutions in specializing general-purpose audio representations for a specific application using speech, a highly demanding field, as an example. We enhance Masked Modeling Duo (M2D), a general-purpose model, to close the performance gap with state-of-the-art (SOTA) speech models. To do so, we propose a new task, denoising distillation, to learn from fine-grained clustered features, and M2D for Speech (M2D-S), which jointly learns the denoising distillation task and M2D masked prediction task. Experimental results show that M2D-S performs comparably to or outperforms SOTA speech models on the SUPERB benchmark, demonstrating that M2D can specialize in a demanding field.

Index Terms: speech representation learning, general-purpose audio representation, denoising, distillation, specialization

1. Introduction

The self-supervised learning general-purpose audio representations (generic models), pre-trained on a large-scale audio dataset, have shown promising performance on a variety of environmental, musical, and speech tasks. While the generic models have demonstrated versatile performance [1, 2, 3, 4, 5], they have not fully demonstrated their usefulness in focused applications such as automatic speech recognition, where SOTA performance is required.

On the other hand, highly advanced speech models have been investigated to meet intense demand. In addition to effective self-supervised learning using contrastive loss and masked prediction [6, 7, 8, 9], SOTA models make use of the learning of discrete representations by quantization [7], pseudo-label generation by clustering [8, 9], and denoising for learning robust representation [9].

We believe generic models can provide even higher performance when pre-trained specifically for a particular application rather than just fine-tuning for the application. This paper explores the challenges and possibilities of specialization using a generic model, masked modeling duo (M2D) [4], in speech as an example. The question we aim to answer is: *Can a general-purpose audio representation be useful in a competitive field?*

To specialize a generic model in speech, we found that the challenges lie in incorporating speech-specific techniques. While a generic model can easily switch the pre-training dataset to a speech corpus, we found it is not yet powerful enough to compete with SOTA speech models. Therefore, resorting to

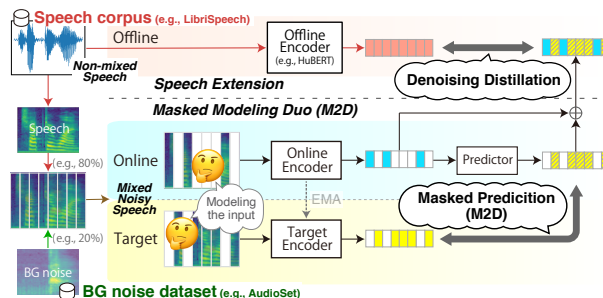


Figure 1: M2D-S extends M2D to speech by adding an offline network for the denoising distillation task. While the online and target networks in M2D learn the masked prediction of the features from noisy speech, the offline network provides pure speech features for learning the denoising distillation.

field-specific extensions (e.g., learning from clustered features) is inevitable.

With a necessary extension, we propose M2D for Speech (M2D-S), which extends M2D with a new task, denoising distillation, to learn from fine-grained clustered features. M2D-S also adapts the pre-training dataset and patch size without changing M2D. Experiments using the SUPERB benchmark [10] show that M2D-S performs comparably to or better than SOTA speech models, demonstrating that a generic model can serve as a pre-training framework for a challenging field.

Our contributions are i) proposal of a denoising distillation task for speech representation learning, ii) proposal of M2D-S to specialize a generic model to speech, and iii) demonstration of the extensibility of a generic model by achieving SOTA performance in speech. Our code is available online¹.

2. Related Work

Self-supervised learning methods that train transformers with masked prediction have shown promising performance in various domains. Speech representation models, such as Mockingjay [11] and TERA [12], take spectrograms as input. TERA employs a masking strategy for splitting both frequency bins and time steps, similar to generic models. SOTA models, such as wav2vec2.0 [7], BigSSL [13], data2vec [14], HuBERT [8], and WavLM [9], typically learn an acoustic feature extraction from the input speech waveform in addition to the representation learning. Notably, models such as wav2vec2.0, HuBERT, and WavLM effectively learn by using discretized pseudo-labels with vector quantization or clustering of pre-trained model features. In addition, WavLM has improved performance on non-

¹<https://github.com/nttcomlab/m2d/tree/master/speech>

ASR tasks through additional masked speech denoising.

Generic models, such as AST [15], SSAST [1], ATST [16], MAE-AST [2], MSM-MAE [3], AudioMAE [17], M2D [4], and BEATs [5], have shown SOTA performance. They typically take spectrograms as acoustic feature input, split input in both the time and frequency axes, and train a Vision Transformer (ViT) [18] while avoiding application-specific techniques such as masking inputs for consecutive time steps, as in speech.

In previous works similar to this paper, Melms et al. [19] and BYOL-S [20] have specialized a generic model BYOL-A [21] in medical and speech applications. SSAST has adapted patch size and pre-training dataset and compared it with speech models on SUPERB. In the NLP domain, LIBERT [22] specializes BERT [23] using an additional task for pre-training a lexically-informed BERT, a similar multi-task learning setting to ours. Previous works that created pseudo-labels or distill models include BEATs, HuBERT, WavLM, DistilHuBERT [24], MT4SSL [25], and RobustDistiller [26]. In particular, concurrent works MT4SSL and RobustDistiller learn multi-tasks, similar to ours, and RobustDistiller auxiliary learns to denoise speech in addition to distillation.

Unlike these previous studies, we attempt to achieve the performance of SOTA speech models by specializing M2D.

3. Method

M2D-S specializes in speech by adding a speech extension network and its task to M2D without changing M2D. Table 1 lists the major differences between M2D and SOTA speech models. The key challenge of speech specialization is to implement in M2D-S the features that critically impact speech task performance.

Table 1: Key differences between models.

Design choices	Speech model (e.g., HuBERT, WavLM)	Generic model (M2D)
(a) Pre-training dataset	Speech (e.g., LibriSpeech [27]) & noises (WavLM)	General audio (e.g., AudioSet [28])
(b) Input data format	Raw waveform	Spectrogram
(c) Input data split	Split in time steps	Split in both freq./time
(d) Feature extractor	CNN	(not used)
(e) Discretized training signals	Clustering pre-trained model features	(not used)
(f) Extra pre-training task	Denosing (WavLM)	(not used)
(g) Masking strategy	Consequent time steps	Random

Based on the preliminary experimental results, we focus on (a) using a speech corpus as the dataset, (c) splitting the input only along the time steps, (e) using clustered pre-trained model features as a training signal, and (f) using a denoising task.

In addition, we propose a new task, denoising distillation, which maximizes the use of fine-grained clusters in the distribution of the pre-trained speech model features and performs (e) and (f) simultaneously.

3.1. Background: Masked Modeling Duo

M2D is a self-supervised learning framework applicable to 2D structured data input such as images and audio spectrograms, and trains ViT with masked prediction. As shown in Fig. 2(a), it consists of two networks, the online and the target, and learns to predict the target output representations using the online output representations. Unlike speech models (e.g., HuBERT), M2D takes a spectrogram (e.g., 80 frequency bins and 208 time steps) as input, which is split into patches (e.g., 16×16) and treated as a series (e.g., $(80/16) \times (208/16) = 65$ patches).

M2D splits the input data x into patches, adds positional encoding, and randomly selects a number of patches according

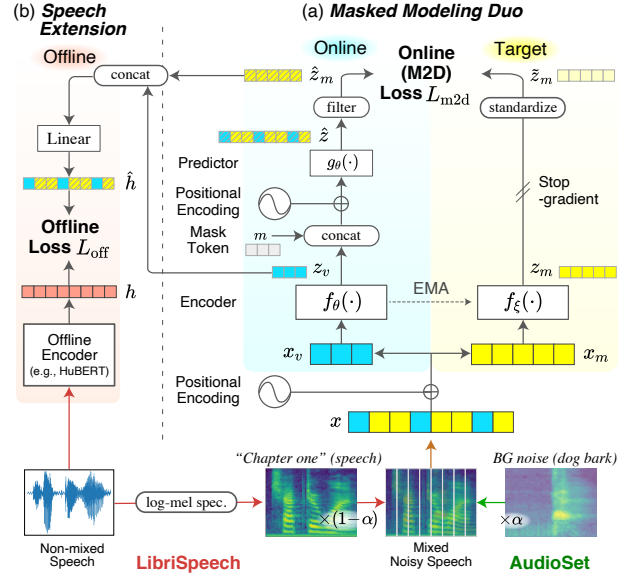


Figure 2: Overview of the M2D-S framework.

to a masking ratio as masked patches x_m (e.g., 60% of the input) and the rest as visible patches x_v (e.g., the remaining 40%).

The online network with a set of weights θ encodes x_v using the online encoder f_θ into the representation $z_v = f_\theta(x_v)$. It concatenates the learnable masked tokens m to z_v , adds the position encoding p , and inputs to the predictor g_θ to predict the representation $\hat{z} = g_\theta(\text{concat}(z_v, m) + p)$. It then outputs the prediction result $\hat{z}_m = \{\hat{z}[i] \mid i \in I_M\}$ of the masked patch representations, where I_M is the set of masked patch indices.

The target network defined by parameter ξ outputs the representation $z_m = f_\xi(x_m)$ and standardizes it to the final target output $\tilde{z}_m = (z_m - \text{mean}(z_m)) / \sqrt{\text{var}(z_m)}$.

The loss is calculated using the online prediction \hat{z}_m against the target output \tilde{z}_m as a training signal by the mean square error (MSE) of l_2 -normalized \hat{z}_m and \tilde{z}_m :

$$L_{m2d} \triangleq \|l_2(\hat{z}_m) - l_2(\tilde{z}_m)\|_2^2 = 2 - 2 \cdot \frac{\langle \hat{z}_m, \tilde{z}_m \rangle}{\|\hat{z}_m\|_2 \cdot \|\tilde{z}_m\|_2}, \quad (1)$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product.

The M2D framework updates parameters θ only by minimizing the loss L_{m2d} as depicted by the stop-gradient in Fig. 2 (a), and updates $\xi \leftarrow \tau\xi + (1 - \tau)\theta$ as an exponential moving average of θ with a decay rate τ .

M2D exploits the momentum encoder to learn effective representations from the target network. After the training, only the f_θ is used as a pre-trained model in downstream tasks.

3.2. Denoising Distillation

We think that a model pre-trained with pseudo labels generated by clustering (e.g., k-means) outputs fine-grained clustered features; therefore, using the pre-trained model features as a training signal should be desirable. HuBERT and WavLM create training signals for the next training iteration by clustering the features produced by the previous iteration model. However, since the previous model has already learned to target clustered labels, we think the model should produce clustered features.

That being said, we propose to combine denoising and distillation as a more flexible approach to learning from clustered targets. The distillation method typically forms a teacher-student network and uses the teacher’s output as a training signal to train the student. Inspired by WavLM, we add noise to the student’s input. In this case, the teacher’s output acts as a

microcluster center corresponding to each data sample, and the student learns to predict cluster centers regardless of the noise. As a result, the denoising distillation is expected to learn from fine-grained clustered features.

3.3. Masked Modeling Duo for Speech

Figure 2 shows M2D-S, which adds an offline network as a speech extension to M2D. We use two types of data (speech and background noise): speech only for the speech extension and noisy speech, a mixture of speech and noise, for M2D.

We distill the offline network as a teacher to M2D as a student. In the speech extension, while the offline encoder produces features h , we concatenate z_v and \hat{z}_m from the output of M2D and project them using a linear layer to predict offline encoder features as \hat{h} . In case M2D splits the input spectrogram in the frequency axis, the \hat{h} is calculated to have one feature prediction per frame by concatenating the feature predictions for each frequency belonging to a time frame.

We calculate the offline loss L_{off} by the MSE of l_2 -normalized h and \hat{h} , the same as Eq. (1) in M2D:

$$L_{\text{off}} \triangleq \|l_2(h) - l_2(\hat{h})\|_2^2 = 2 - 2 \cdot \frac{\langle h, \hat{h} \rangle}{\|h\|_2 \cdot \|\hat{h}\|_2} \quad (2)$$

The overall M2D-S loss L_{m2dS} is then calculated by combining L_{m2d} and L_{off} :

$$L_{\text{m2dS}} = \lambda_{\text{m2d}} L_{\text{m2d}} + \lambda_{\text{off}} L_{\text{off}}, \quad (3)$$

where the loss weights λ_{m2d} and λ_{off} control the contribution.

The noisy speech is a mixture of background noise and speech in the α to $(1 - \alpha)$ ratio, where α is a dataset noise ratio.

4. Experiments

First, we validate the effect of each modification made in M2D-S in ablation studies of the pre-training dataset (Section 4.2), the input size (Section 4.3), and the pre-training tasks (Section 4.4). Finally, we compare M2D-S combining all the best practices with SOTA models (Section 4.5).

4.1. Experimental Setup

We used the same M2D configurations in M2D-S, including the use of ViT Base as the encoder, with a fixed masking ratio of 0.6. For each experiment, the parameters of interest were varied with the following defaults:² dataset noise ratio α of 0.2, input duration T of 2.08 s, and patch size of 80×4 . M2D takes log-mel spectrogram features as input, which we preprocessed with a sampling frequency of 16,000 Hz, window size of 25 ms, hop size of 10 ms, and frequency bins F of 80 in the range of 50 to 8,000 Hz and normalized using dataset statistics.

Pre-training details We set the number of epochs to 1,000 and the warm-up epochs to 60, where a single epoch consumes all LibriSpeech samples. All other settings were the same as in the M2D, including the batch size and optimizer settings. We used the 9th transformer layer output of the 2nd-iteration HuBERT Base model³ as an offline encoder as in WavLM.

For the speech corpus, we used LibriSpeech with 281,241 samples (960 h) from all training splits (LS-960). For the background noise dataset, we used AudioSet [28] (AS) with

²We provide complete details at: <https://github.com/nttcslab/m2d/tree/master/speech>

³<https://huggingface.co/facebook/hubert-base-1s960>

2,005,132 samples (5,569 h) of 10-s audio from the balanced and unbalanced train segments. We randomly sampled then randomly cropped a segment for the input duration and mixed it with the speech. When mixing, the log-mel spectrograms were once reverted to a linear scale, mixed according to the ratio, and then converted back to log scale.

Evaluation details We evaluated all pre-trained models in SUPERB [10] for the speech task performance and a linear evaluation for the non-speech task. In the SUPERB evaluation, the model weights were frozen, and the weighted sum of the features from all transformer layers was used in the evaluation. The evaluation tasks include phoneme recognition (PR), keyword spotting (KS), intent classification (IC), speaker identification (SID), and emotion recognition (ER).

We used linear evaluation as a supplemental measure to assess non-speech task performance when specializing in speech. The linear evaluation and task details are the same as in BYOL-A [21]. We report the average accuracies of two environmental sound tasks (ENV) and three music tasks (MUS). In this evaluation, we trained a linear layer using only the final layer feature of the frozen model instead of using all layer features as in SUPERB.

4.2. Pre-training Dataset Ablations

We assessed the impact of the dataset used for pre-training by comparing the performance when using the target domain dataset, i.e., LibriSpeech, and the general audio dataset, i.e., AudioSet. To do so, we used M2D without speech extension and varied the dataset noise ratio α . When only AudioSet was used, we defined one epoch as the number of LibriSpeech samples, and we randomly sampled from AudioSet.

The results in Table 2 show that the speech tasks perform best when only using LibriSpeech, and performance deteriorates as the ratio of AudioSet increases. Conversely, the non-speech tasks (ENV, MUS) perform best when only using AudioSet, and performance deteriorates as the ratio of LibriSpeech increases. Exceptionally, SID performed best when the ratio $\alpha = 0.2$ and there was background noise in the speech. However, the use of a speech domain dataset is generally considered to contribute to speech task performance.

Table 2: M2D pre-training dataset noise ratio ablations. (input duration $T = 2.08s$ and patch size 80×4)

Dataset noise ratio α	PR	KS	IC	SID	ER	ENV	MUS
	PER↓	Acc↑	Acc↑	Acc↑	Acc↑	Acc↑	Acc↑
0.0 (LS-960 only)	10.98	96.20	95.02	74.77	63.02	66.54	49.83
0.1	11.97	91.54	94.38	77.89	61.44	72.77	51.68
0.2	11.87	93.30	93.28	78.46	61.75	73.57	53.76
0.3	12.19	93.38	94.46	78.25	61.55	74.67	54.44
0.4	12.39	95.15	94.15	77.15	61.52	75.16	54.39
0.5	12.53	93.70	92.14	76.58	61.07	76.38	54.73
1.0 (AudioSet only)	27.04	81.91	82.78	68.54	60.85	83.31	61.88

4.3. Input Size Ablations

The patch splitting of the input spectrogram in the generic model is a key difference from the speech model. In addition to patch size, we examined the impact on performance of different input durations using M2D without speech extension.

Table 3 patch size results indicate that the setting equivalent to the speech model (80×2) has the best performance balance except for SID; for patch size 80×2 (Frequency \times Time-step), 80 indicates no splitting on the frequency axis, and 2 corresponds to 20 ms per frame. The results also show that the longer time steps (lower frame rates) degrade the performance, especially in PR, similar to the results in [29].

The patch sizes of 40×4 and 40×2 , which assign two patches along the frequency axis, showed the best results for PR and SID, but inferior results on some other tasks. Overall, these results indicate that no patch split along the frequency axis provides a balanced better performance on the speech tasks. Notably, our results align with SSAST [1].

The input duration results in Table 4 show a trend that the longer the input, the better the results. The trend of better performance when learning from features in longer series is consistent as it is in M2D.

Table 3: *M2D patch size ablations.*
(dataset noise ratio $\alpha = 0.2$ and input duration $T = 2.08s$)

Patch size Freq. \times Time	PR	KS	IC	SID	ER	ENV	MUS
	PER \downarrow	Acc \uparrow	Acc \uparrow	Acc \uparrow	Acc \uparrow	Acc \uparrow	Acc \uparrow
16×16 (M2D [4])	77.92	86.24	83.23	79.65	58.88	80.27	59.60
40×8	29.28	90.25	89.67	77.51	59.39	78.56	56.62
40×4	11.49	91.77	90.93	81.64	60.30	80.00	57.48
40×2	15.14	81.47	82.73	85.44	61.49	80.98	58.67
80×8	30.28	89.69	90.51	75.66	58.97	72.62	54.28
80×4	11.87	93.30	93.28	78.46	61.75	73.57	53.76
80×2 (\Leftrightarrow speech models)	11.74	93.28	93.51	78.86	60.67	74.98	53.28

Table 4: *M2D input duration ablations.*
(dataset noise ratio $\alpha = 0.2$ and patch size 80×4)

Input duration T	PR	KS	IC	SID	ER	ENV	MUS
	PER \downarrow	Acc \uparrow	Acc \uparrow	Acc \uparrow	Acc \uparrow	Acc \uparrow	Acc \uparrow
$T=2.08s$	11.87	93.30	93.28	78.46	61.75	73.57	53.76
$T=3.04s$	9.81	94.99	95.15	79.18	63.39	74.36	52.14
$T=4.00s$	8.50	95.04	94.83	81.24	63.81	74.11	52.23
$T=5.12s$	8.10	95.24	94.70	78.73	65.47	74.69	51.66
$T=6.08s$	7.74	96.41	95.50	80.48	64.06	72.42	50.64

Table 5: *M2D-S pre-training task ablations.*
(input duration $T = 2.08s$, and patch size 80×4)

	Offline		Online	PR	KS	IC	SID	ER	ENV	MUS
	denoise	distill.	M2D	PER \downarrow	Acc \uparrow	Acc \uparrow	Acc \uparrow	Acc \uparrow	Acc \uparrow	Acc \uparrow
(a)			✓	11.87	93.30	93.28	78.46	61.75	73.57	53.76
(b)	✓			8.10	94.33	94.81	65.51	61.27	56.96	44.16
(c)	✓	✓		7.80	93.88	93.38	70.03	61.58	42.42	35.36
(d)		✓	✓	8.78	90.90	91.35	62.68	58.63	57.15	43.75
(e)	✓	✓	✓	7.02	97.36	97.10	78.12	64.09	55.49	41.51

4.4. Pre-training Task Ablations

We assessed the effectiveness of the denoising distillation and the M2D task and also tested without denoising. We switched the tasks in the offline and M2D networks by setting λ_{off} and λ_{m2d} in Eq. (3) to 0 or 1.0, respectively, and switched the denoising task by setting α to 0 (disabled) or 0.2 (enabled).

The results in Table 5 show that (b) distillation only is better than (a) M2D or (d) M2D plus distillation, indicating that the distillation of a speech model trained with clustered features is effective for the speech tasks. Compared to (b) with only distillation, (c) with additional denoising performs better except for KS and IC, indicating the effectiveness of the denoising task. M2D-S configuration (e), learning all tasks together, significantly improves task performance, showing that combining denoising distillation and M2D is remarkably effective.

4.5. Comparison with SOTA

While previous experiments have improved performance, they have not reached SOTA performance; therefore, we combine all the settings that offered the best performance in M2D-S. Table 6 compares the original M2D, proposed M2D-S, and the SOTA speech *Base* models, which have a number of parameters close to that in M2D. We tested three input durations T of M2D-S while keeping the other settings fixed. We also optimized λ_{off} and λ_{m2d} by parameter search.

Table 6: *Comparison with SOTA speech models.*
($\lambda_{\text{off}} = 0.5$, $\lambda_{\text{m2d}} = 1$, $\alpha = 0.2$, and patch size 80×2)

Model	Dataset	PR	KS	IC	SID	ER	ENV	MUS
		PER \downarrow	Acc \uparrow	Acc \uparrow	Acc \uparrow	Acc \uparrow	Acc \uparrow	Acc \uparrow
wav2vec2.0 Base [7] [†]	LS-960	5.74	96.23	92.35	75.18	63.43	37.66	32.02
HuBERT Base [8] [†]	LS-960	5.41	96.30	98.34	81.42	64.92	62.76	46.26
WavLM Base [9] [†]	LS-960+DNS	4.84	96.79	98.63	84.51	65.94	54.45	40.98
<i>(Proposed using ViT Base)</i>								
M2D-S $T=4.0s$	LS-960+AS	5.72	97.80	97.80	81.97	66.36	53.22	41.71
M2D-S $T=5.12s$	LS-960+AS	5.64	97.76	97.65	80.69	65.35	57.34	43.23
M2D-S $T=6.08s$	LS-960+AS	5.33	97.81	97.63	81.74	66.13	54.77	43.75
<i>(Conventional using ViT Base)</i>								
M2D ratio=0.6 [4] [‡]	AS	78.30	80.23	76.77	80.68	61.17	88.63	66.56
<i>(Reference Large models)</i>								
wav2vec2.0 Large [7] [†]	LL-60k	4.75	96.66	95.28	86.14	65.64	60.82	42.75
HuBERT Large [8] [†]	LL-60k	3.53	95.29	98.76	90.33	67.62	59.51	44.35
WavLM Large [9] [†]	Mix-94k	3.06	97.86	99.31	95.49	70.62	69.32	50.56

[†] ENV and MUS results were obtained using publicly available pre-trained models.

[‡] The original M2D takes input with $T = 6.08s$ and uses a patch size of 16×16 .

The results show that M2D-S significantly improves the performance of M2D to a level comparable to SOTA speech models. Furthermore, M2D-S outperforms these SOTA models on KS and ER tasks. Notably, this is achieved without a CNN acoustic feature extractor used by the previous methods, while M2D-S takes log-mel spectrograms as input. On the other hand, M2D-S performed worse on the non-speech ENV and MUS tasks than M2D and HuBERT, which is interestingly similar to WavLM. Both WavLM and M2D-S learn speech denoising, which we suspect may be the cause of this side effect.

5. Discussion for Future Specialization

This study explored the requirements for specializing M2D in speech, and the lessons learned may be useful for specialization in other acoustic fields.

- Optimizing the patch size may be effective for other applications. While splitting only the time axis was effective for speech, splitting along both the frequency and time axes has been reported as effective for other audio tasks [1, 3].
- In addition to using a dataset from the target domain, combining a denoising task with background noise may be effective in learning application-focused representations.
- Learning the denoising distillation task requires an appropriate teacher model. Although we used a model that learned pseudo-labels by clustering, pre-trained classifier models available in the application field may be effective since they learn from class labels, considered human-curated clusters.

6. Conclusion

This paper explored the challenges and solutions in specializing a general-purpose audio representation (generic model) to a specific domain using speech as an example. We found the challenge is to incorporate speech-specific techniques, such as learning from clustered features.

To achieve speech SOTA performance, we proposed a new task, denoising distillation, to learn from fine-grained clustered features, and M2D for Speech (M2D-S), which jointly learns the denoising distillation task and M2D masked prediction task. In the experiments on the SUPERB benchmark, M2D-S significantly improved the performance of M2D. It performed comparable to SOTA speech models (e.g., HuBERT and WavLM) and outperformed them in keyword spotting and emotion recognition tasks, demonstrating that a generic model can specialize and be useful in a challenging field.

Our study indicated that a generic model is potentially useful for specializing in other audio applications. Our code is available online for future studies.

7. References

- [1] Y. Gong, C.-I. Lai, Y.-A. Chung, and J. Glass, “SSAST: Self-Supervised Audio Spectrogram Transformer,” in *AAAI*, vol. 36, no. 10, 2022, pp. 10 699–10 709.
- [2] A. Baade, P. Peng, and D. Harwath, “MAE-AST: Masked Autoencoding Audio Spectrogram Transformer,” in *Interspeech*, 2022, pp. 2438–2442.
- [3] D. Niizumi, D. Takeuchi, Y. Ohishi, N. Harada, and K. Kashino, “Masked Spectrogram Modeling using Masked Autoencoders for Learning General-purpose Audio Representation,” in *HEAR: Holistic Evaluation of Audio Representations (NeurIPS 2021 Competition)*, vol. 166, 2022, pp. 1–24.
- [4] —, “Masked Modeling Duo: Learning Representations by Encouraging Both Networks to Model the Input,” in *ICASSP*, 2023, pp. 1–5.
- [5] S. Chen, Y. Wu, C. Wang, S. Liu, D. Tompkins, Z. Chen, and F. Wei, “BEATs: Audio Pre-Training with Acoustic Tokenizers,” *arXiv preprint arXiv:2212.09058*, 2022.
- [6] S. Liu, A. Mallol-Ragolta, E. Parada-Cabaleiro, K. Qian, X. Jing, A. Kathan, B. Hu, and B. W. Schuller, “Audio self-supervised learning: A survey,” *Patterns*, vol. 3, no. 12, p. 100616, 2022.
- [7] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *NeurIPS*, 2020.
- [8] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, “HuBERT: Self-Supervised Speech Representation Learning by Masked Prediction of Hidden Units,” *IEEE/ACM Trans. Audio, Speech, Language Process.*, p. 3451–3460, 2021.
- [9] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, J. Wu, M. Zeng, X. Yu, and F. Wei, “WavLM: Large-Scale Self-Supervised Pre-Training for Full Stack Speech Processing,” *IEEE J. Sel. Top. Signal Process.*, vol. 16, no. 6, p. 1505–1518, 2022.
- [10] S. wen Yang, P.-H. Chi, Y.-S. Chuang, C.-I. J. Lai, K. Lakhota, Y. Y. Lin, A. T. Liu, J. Shi, X. Chang, G.-T. Lin, T.-H. Huang, W.-C. Tseng, K. tik Lee, D.-R. Liu, Z. Huang, S. Dong, S.-W. Li, S. Watanabe, A. Mohamed, and H. yi Lee, “SUPERB: Speech Processing Universal PERFORMANCE Benchmark,” in *Interspeech*, 2021, pp. 1194–1198.
- [11] A. T. Liu, S.-w. Yang, P.-H. Chi, P.-c. Hsu, and H.-y. Lee, “Mockingjay: Unsupervised Speech Representation Learning with Deep Bidirectional Transformer Encoders,” in *ICASSP*, 2020, pp. 6419–6423.
- [12] A. T. Liu, S.-W. Li, and H.-y. Lee, “TERA: Self-Supervised Learning of Transformer Encoder Representation for Speech,” *IEEE/ACM Trans. Audio, Speech, Language Process.*, vol. 29, pp. 2351–2366, 2021.
- [13] Y. Zhang, D. S. Park, W. Han, J. Qin, A. Gulati, J. Shor, A. Jansen, Y. Xu, Y. Huang, S. Wang, Z. Zhou, B. Li, M. Ma, W. Chan, J. Yu, Y. Wang, L. Cao, K. C. Sim, B. Ramabhadran, T. N. Sainath, F. Beaufays, Z. Chen, Q. V. Le, C.-C. Chiu, R. Pang, and Y. Wu, “BigSSL: Exploring the frontier of large-scale semi-supervised learning for automatic speech recognition,” *IEEE J. Sel. Top. Signal Process.*, vol. 16, no. 6, p. 1519–1532, 2022.
- [14] A. Baevski, W.-N. Hsu, Q. Xu, A. Babu, J. Gu, and M. Auli, “data2vec: A general framework for self-supervised learning in speech, vision and language,” in *ICML*, 2022, pp. 1298–1312.
- [15] Y. Gong, Y.-A. Chung, and J. Glass, “AST: Audio Spectrogram Transformer,” in *Interspeech*, 2021, pp. 571–575.
- [16] X. LI and X. Li, “ATST: Audio Representation Learning with Teacher-Student Transformer,” in *Interspeech*, 2022, pp. 4172–4176.
- [17] P.-Y. Huang, H. Xu, J. Li, A. Baevski, M. Auli, W. Galuba, F. Metze, and C. Feichtenhofer, “Masked autoencoders that listen,” in *NeurIPS*, 2022.
- [18] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, “An image is worth 16x16 words: Transformers for image recognition at scale,” in *ICLR*, 2021.
- [19] L. Melms, R. R. Ilesan, U. Köhler, O. Hildebrandt, R. Conradt, J. Eckstein, C. Atila, S. Matrood, B. Schieffer, J. R. Schaefer, T. Müller, J. Obergassel, N. Schlicker, and M. C. Hirsch, “Training one model to detect heart and lung sound events from single point auscultations,” *arXiv preprint arXiv:2301.06078*, 2023.
- [20] G. Elbanna, N. Scheidwasser-Clow, M. Kegler, P. Beckmann, K. El Hajal, and M. Cernak, “BYOL-S: Learning self-supervised speech representations by bootstrapping,” in *HEAR: Holistic Evaluation of Audio Representations (NeurIPS 2021 Competition)*, vol. 166, 2022, pp. 25–47.
- [21] D. Niizumi, D. Takeuchi, Y. Ohishi, N. Harada, and K. Kashino, “BYOL for Audio: Exploring pre-trained general-purpose audio representations,” *IEEE/ACM Trans. Audio, Speech, Language Process.*, vol. 31, p. 137–151, 2023.
- [22] A. Lauscher, I. Vulić, E. M. Ponti, A. Korhonen, and G. Glavaš, “Specializing unsupervised pretraining models for word-level semantic similarity,” in *COLING*, 2020, pp. 1371–1383.
- [23] J. Devlin, M. Chang, K. Lee, and K. Toutanova, “BERT: pre-training of deep bidirectional transformers for language understanding,” in *NAACL-HLT*, 2019, pp. 4171–4186.
- [24] H.-J. Chang, S.-w. Yang, and H.-y. Lee, “DistilHuBERT: Speech representation learning by layer-wise distillation of hidden-unit bert,” in *ICASSP*, 2022, pp. 7087–7091.
- [25] Z. Ma, Z. Zhen, C. Tang, Y. Wang, and X. Chen, “MT4SSL: Boosting Self-Supervised Speech Representation Learning by Integrating Multiple Targets,” *to appear at Interspeech*, 2023.
- [26] H. R. Guimarães, A. Pimentel, A. R. Avila, M. Rezagholizadeh, B. Chen, and T. H. Falk, “RobustDistiller: Compressing Universal Speech Representations for Enhanced Environment Robustness,” in *ICASSP*, 2023, pp. 1–5.
- [27] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An asr corpus based on public domain audio books,” in *ICASSP*, 2015, pp. 5206–5210.
- [28] J. F. Gemmeke, D. P. W. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, “Audio Set: An ontology and human-labeled dataset for audio events,” in *ICASSP*, 2017, pp. 776–780.
- [29] Y. Meng, H.-J. Chen, J. Shi, S. Watanabe, P. Garcia, H.-y. Lee, and H. Tang, “On compressing sequences for self-supervised speech models,” in *SLT*, 2022, pp. 1128–1135.