



# Downstream Task Agnostic Speech Enhancement with Self-Supervised Representation Loss

Hiroshi Sato, Ryo Masumura, Tsubasa Ochiai, Marc Delcroix,  
Takafumi Moriya, Takanori Ashihara, Kentaro Shinayama,  
Saki Mizuno, Mana Ihori, Tomohiro Tanaka, Nobukatsu Hojo

NTT Corporation, Japan

hiroshi.satou.bh@hco.ntt.co.jp

## Abstract

Self-supervised learning (SSL) is the latest breakthrough in speech processing, especially for label-scarce downstream tasks by leveraging massive unlabeled audio data. The noise robustness of the SSL is one of the important challenges to expanding its application. We can use speech enhancement (SE) to tackle this issue. However, the mismatch between the SE model and SSL models potentially limits its effect. In this work, we propose a new SE training criterion that minimizes the distance between clean and enhanced signals in the feature representation of the SSL model to alleviate the mismatch. We expect that the loss in the SSL domain could guide SE training to preserve or enhance various levels of characteristics of the speech signals that may be required for high-level downstream tasks. Experiments show that our proposal improves the performance of an SE and SSL pipeline on five downstream tasks with noisy input while maintaining the SE performance.

**Index Terms:** speech enhancement, self-supervised learning, SUPERB benchmark, denoising, deep learning

## 1. Introduction

In the past decade, advances in deep learning have drastically improved speech processing technology and its real-world implementation. The supervised learning framework has been a central component of these advancements, as it improves the performance in proportion to the amount of available labeled data [1, 2]. However, the annotation cost of labeled data constrains performance, particularly in tasks where a large amount of labeled data is either unavailable or too expensive to obtain.

To address this issue, self-supervised learning (SSL), a machine learning technique that can learn useful representations without the need for artificially labeled data, is attracting research attention [3–5]. One research goal of SSL in speech processing is to obtain a single SSL model whose task-agnostic representations are effective for various downstream tasks [6–10]. To this end, various techniques have been proposed, among which wav2vec 2.0 [11], HuBERT [12], and WavLM [13] are some of the most widely applied approaches. Indeed they provide effective representations for various downstream tasks [14, 15].

The noise robustness of the SSL is one of the critical challenges to expanding its application, although this issue has not yet been fully examined. Among the leading SSL models, WavLM [13] addresses the issue of robustness against noise and overlapping speech from interfering background speakers by incorporating a denoising framework into the pre-training process. The prior study showed that it was relatively robust to noise [16]. Besides, Chung *et al* revealed that the noise robustness of SSL models could be further improved by introducing a single-channel speech enhancement (SE) frontend [17].

They showed that the SE, SSL, and automatic speech recognition (ASR) pipeline achieves state-of-the-art performance on the monaural CHiME-4 task. However, they also reported that the fine-tuning of SE frontend on the downstream ASR task is critically necessary to benefit from the SE frontend. In other words, the SE model trained separately from SSL models with the SE training objective suffers from a mismatch with SSL models, potentially limiting SE performance.

To address this issue, we propose an SSL Mean Square Error (SSL-MSE) loss as way of learning an SE model optimal as the frontend of the SSL model while being independent of the downstream tasks, e.g., ASR. Contrary to the conventional SE training objective that minimizes the distance between enhanced speech and clean source in the time or spectrum domain, SSL-MSE aims to bring these signals closer in the feature domain projected by the SSL model, which makes enhanced signal well suited for the SSL model. Moreover, we propose using a multi-task training objective that combines the SSL-MSE loss with a conventional SE loss to preserve SE capabilities of the front-end. Since SSL models have been shown to learn effective representations for many downstream tasks, they can capture not only acoustic but also higher-level information such as phonetic or semantic information [18–20]. Thus, we expect that creating a loss term in the SSL domain could guide SE training to preserve or enhance various levels of characteristics of the speech signals that may be required for high-level downstream tasks.

A major difference between SSL-MSE training and the fine-tuning approach of the SE model with a downstream task adopted in [17] is that ours can train the SE model in a downstream task agnostic manner. In [17], the SE frontend and downstream ASR models are jointly trained on the ASR criterion while freezing the SSL model. On the other hand, our proposal constructs a single SE model suitable for a general SSL pipeline that can be applied to various downstream tasks without requiring downstream task-specific fine-tuning. This allows sharing a single SE and SSL model for multiple downstream tasks by only training the relatively small-sized downstream model for each task.

We confirmed the effectiveness of the proposed SSL-MSE loss with various downstream tasks by performing experiments on a noisy version of a subset of the SUPERB benchmark [14]. The result shows that the proposed SSL-MSE substantially improves the performance of downstream tasks using WavLM. Our proposal is also effective even with noise-robust downstream models trained with noisy paired data. Moreover, when using the multitask loss, the proposed SE system maintains SE performance, allowing using it also for applications requiring access to the enhanced speech signal, e.g., listening applications. It can also contribute to the interpretability of the process. An interesting finding of our study is that SE systems trained

with SSL-MSE loss generalize to downstream tasks using different SSL models from those used for the SSL-MSE training. This seems to confirm our intuition that the task-agnostic feature representation of SSL models helps us to learn a more robust SE frontend.

## 2. Conventional Method

### 2.1. Task-agnostic SSL upstream models

Various types of SSL methods have been proposed to extract powerful features or representations from speech recordings that are useful for downstream tasks. We can write the feature extraction process with the SSL model as follows:

$$\mathbf{F}_{1:N} = \text{SSL}(\mathbf{x}; \theta_{\text{ssl}}), \quad (1)$$

where  $\mathbf{x} \in \mathbb{R}^T$  denotes the monaural input signal in raw waveform with  $T$  samples,  $\mathbf{F}_n \in \mathbb{R}^{D \times T'}$  denotes the time series of extracted features obtained from the  $n$ -th layer of the SSL model,  $N$  is the number of layers of the SSL model, and  $\theta_{\text{ssl}}$  denotes the learnable parameters of the SSL model.  $T'$  and  $D$  denote the number of frames and dimensions of the extracted features, respectively. The learnable parameters,  $\theta_{\text{ssl}}$  are trained using pretext tasks with unpaired audio-only data. For example, HuBERT [12] and WavLM [13] are trained using the BERT-like masked prediction task [21] on the target label generated in an offline clustering step. Of particular note, WavLM is made robust to noise and interference speakers through the addition of Deep Noise Suppression (DNS) noise [22] and interference speakers during its training. Various types of pre-trained SSL models have been made publicly available for application to downstream tasks.

### 2.2. Task specific downstream models

There are two major ways to apply SSL models to downstream tasks, either as a fixed feature extractor, or permitting retraining of their parameters [3]. Following previous studies [14, 16, 17], we adopted the former approach, i.e., freezing the parameters of the SSL model  $\theta_{\text{ssl}}$  during the training of the downstream models, because SSL models are usually very large and thus, it is computationally too intensive to fine-tune the SSL models for every downstream task. The downstream task-specific additional layers are relatively small deep neural networks, which we call the downstream model.

An effective way to apply the SSL features across various types of downstream tasks is to use the weighted sum of the embeddings from different layers in SSL model as the input feature of the downstream model [13, 14]. The process of the downstream model for a task  $\tau$  can be written as follows:

$$\mathbf{F}^\tau = \text{WS}(\mathbf{F}_{1:N}; \theta_w^\tau) \equiv \sum_{n=1}^N w_n^\tau \mathbf{F}_n, \quad (2)$$

$$\hat{\mathbf{I}}^\tau = \text{DOWNSTREAM}^\tau(\mathbf{F}^\tau; \theta_{\text{ds}}^\tau), \quad (3)$$

where  $\text{WS}(\cdot)$  represents the weighted sum function,  $\theta_w^\tau = [w_1^\tau, \dots, w_N^\tau]$  are the weights for latent representations obtained from each layer, which are learnable parameters, and  $\hat{\mathbf{I}}^\tau$  is the estimation result attained by the downstream model.  $\text{DOWNSTREAM}^\tau(\cdot)$  denotes the downstream model for the task  $\tau$  whose learnable parameters are  $\theta_{\text{ds}}^\tau$ . The learnable parameters ( $\theta_w^\tau, \theta_{\text{ds}}^\tau$ ) are jointly optimized by using task-specific paired data ( $\mathbf{I}^\tau, \hat{\mathbf{I}}^\tau$ ), while the upstream model  $\theta_{\text{ssl}}$  is frozen.




Figure 1: Overview of the proposed SSL-MSE loss.

### 2.3. Speech Enhancement and SSL pipeline

Speech enhancement is a method to suppress undesired audio from the input noisy observations. We denote the SE process as follows:

$$\hat{\mathbf{x}} = \text{SE}(\mathbf{y}; \theta_{\text{se}}), \quad (4)$$

where  $\mathbf{y} \in \mathbb{R}^T$  and  $\hat{\mathbf{x}} \in \mathbb{R}^T$  denote the noisy observations and enhanced signals, respectively.  $\theta_{\text{se}}$  denotes the learnable parameters of the SE model, which is optimized by minimizing the distance between enhanced signal  $\hat{\mathbf{x}}$  and the ground truth clean source  $\mathbf{x} \in \mathbb{R}^T$ . Specifically, we adopt the scale-dependent signal-to-noise ratio loss  $\mathcal{L}_{\text{SNR}}(\cdot)$  as the distance measure, which is defined as follows:

$$\mathcal{L}_{\text{SNR}} = 10 \log_{10} \frac{\|\hat{\mathbf{x}}\|^2}{\|\mathbf{x} - \hat{\mathbf{x}}\|^2}. \quad (5)$$

Since the clean source is not available for real-recorded noisy speech, simulated mixtures of clean source  $\mathbf{x}$  and noise are commonly used as noisy speech  $\mathbf{y}$  in SE training. The entire SE and SSL pipeline for downstream task  $\tau$  can be written as follows:

$$\hat{\mathbf{I}}^\tau = \text{DOWNSTREAM}^\tau(\text{WS}(\text{SSL}(\text{SE}(\mathbf{y}; \theta_{\text{se}}); \theta_{\text{ssl}}); \theta_w^\tau); \theta_{\text{ds}}^\tau). \quad (6)$$

It is commonly known that the mismatch between SE frontend and backends can degrade the performance of backend tasks when the SE frontend is trained separately from backends, using the SE training criterion [23–25]. For the SE, SSL, and ASR pipeline, this mismatch was mitigated by using the downstream ASR loss in jointly retraining SE and downstream models [16, 17]. However, this method may not be feasible when considering application for more than a few downstream tasks, because the retraining of the SE model for every downstream task is computationally demanding as the SE model tend to have many more learnable parameters than those of the downstream models to maintain the enhancement performance. Moreover, the fine-tuning on the downstream task makes the SE frontend task-dependent, which makes it infeasible to share a single general SE model for multiple downstream tasks.

## 3. Proposed Method

To obtain the downstream task-agnostic SE frontend while reducing the mismatch, we propose SSL-MSE. Figure 1 shows the general framework of the proposed SSL-MSE training. In addition to the conventional signal-to-noise ratio (SNR) loss, we calculate SSL-MSE which drives the enhanced signal closer to the clean source in the SSL feature domain. Formally, SSL-MSE  $\mathcal{L}_{\text{SSL-MSE}}$  between enhanced speech  $\hat{\mathbf{x}} = \text{SE}(\mathbf{y}; \theta_{\text{se}})$  and clean source  $\mathbf{x}$  is calculated as the mean square error between SSL features extracted from these signals as follows:

$$\mathcal{L}_{\text{SSL-MSE}} = \|\bar{\mathbf{F}}^{\text{ench}} - \bar{\mathbf{F}}^{\text{clean}}\|_{\text{F}}^2 / DT', \quad (7)$$

$$\bar{\mathbf{F}}^{\text{ench}} = \sum_{n=1}^N \tilde{w}_n \mathbf{F}_n^{\text{ench}}, \quad \bar{\mathbf{F}}^{\text{clean}} = \sum_{n=1}^N \tilde{w}_n \mathbf{F}_n^{\text{clean}}, \quad (8)$$

$$\mathbf{F}_{1:N}^{\text{ench}} = \text{SSL}(\text{SE}(\mathbf{y}; \theta_{\text{se}}); \theta_{\text{ssl}}), \quad \mathbf{F}_{1:N}^{\text{clean}} = \text{SSL}(\mathbf{x}; \theta_{\text{ssl}}), \quad (9)$$

where  $\|\cdot\|_F$  represents Frobenius norm and  $\tilde{\mathbf{w}} = [\tilde{w}_1, \dots, \tilde{w}_N]$  represents the layer weight for the latent feature obtained from each layer of the SSL model. The layer weight is a hyperparameter that controls which layer output SSL-MSE focuses on. We tested three types of weighting: 1) *last* uses only the last layer output, 2) *all* weighs every layer equally, and 3) *latter-half* weights uniformly over the latter half layers. The model parameter  $\theta_{se}$  is optimized by minimizing the multitask loss function  $\mathcal{L}$  expressed as follow:

$$\mathcal{L} = \mathcal{L}_{SSL-MSE} + \alpha \mathcal{L}_{SNR}, \quad (10)$$

where  $\alpha$  denotes the multitask weight for the SNR loss [26]. We hereafter refer to the hyperparameter  $\alpha$  as SNR weight. Note that the parameters of the SSL model  $\theta_{ssl}$  are frozen while the SSL-MSE training.

We adopted WavLM [13] as a SSL model to calculate SSL-MSE loss. Note that SSL-MSE training does not depend on the downstream tasks, only on the SSL model. How the SSL-MSE training generalizes over other types of SSL models than that used for SSL-MSE training is discussed in Section 4.2.1.

## 4. Experiments

### 4.1. Experimental setup

#### 4.1.1. Speech enhancement frontend

The SE model was trained on simulated mixtures of speech and noise data. We used LibriSpeech [27] for speech recordings and DNS noise [22] for noise recordings. The number of noisy observations were 100,000 and 5,000 for the training and development sets, respectively. The noise was added at SNR values randomly sampled from -3 to 20 dB. We chose DNS noise to create a fair comparison of the SSL pipeline with/without the SE module, as DNS noise is also used in training the WavLM models. We adopted ConvTasnet as the SE front-end module, which converts noisy raw-waveform audio into enhanced raw-waveforms in a time-domain, end-to-end processing manner [28]. According to the setup adopted in [29], we set the hyperparameters to  $N=4096$ ,  $L=320$ ,  $B=256$ ,  $R=4$ ,  $X=8$ ,  $H=512$  and  $P=3$  following the notation in [28]. The SSL-MSE loss was introduced after pre-training the model with conventional SNR loss to speed up the conversion. The initial learning rate for the pre-training with SNR loss was set to  $5e-4$  and was multiplied by  $3/4$  if the loss on the development set did not decrease for 2 epochs. For optimization, we adopted the Adam optimizer [30]. The models were pre-trained for 100 epochs.

#### 4.1.2. SSL-MSE loss

To calculate SSL-MSE, we used off-the-shelf WavLM BASE+ model. We tested three configuration of layer weight  $\tilde{\mathbf{w}}$ : last, all, and latter-half as defined in Section 3. We used the last setup unless otherwise specified. We tested the SNR weight  $\alpha$  within  $\{0, 0.0001, 0.001, 0.01, 0.1, 1\}$ . We adopted 0.1 unless otherwise specified, which is the smallest value where SNR loss on development set did not degrade compared with that of conventional SNR loss training. We fine-tuned the SE model using SSL-MSE loss with the initial learning rate of  $1e-4$  and for up to 50 epochs.

#### 4.1.3. SSL and downstream model

We implemented the SSL pipeline using the S3PRL toolkit [14, 31]. As the SSL model, we used the publicly available pre-trained models of wav2vec 2.0 BASE, Hubert BASE, WavLM BASE+, and WavLM LARGE without updating their param-

eters. We tested two setups for downstream model training: 1) *official* setup where downstream models were trained with SUPERB official training data that is relatively ‘clean’ speech data, according to S3PRL SUPERB recipe [31], and 2) *noise robust* setup where DNS noise was added to SUPERB official downstream training data at SNR values randomly sampled from -3 to 20 dB. The results gained in the noise robust setup are shown in Table 3; for the other experiments we adopt the official setup. SE frontend was not applied while training the downstream model in noise robust setup as well as in official setup. The downstream models were prepared for each downstream task for each SSL model.

#### 4.1.4. Evaluation details

In this work, we evaluate the SE model in terms of 1) speech enhancement performance and 2) the performance of downstream tasks. SE performance was measured using the 3,000 simulated noisy observations of LibriSpeech data and DNS noise at SNR values randomly sampled from 0 to 10 dB. As the performance measure, we adopted average value of the scale-independent source-to-distortion ratio (SDR) [32] and perceptual evaluation of speech quality (PESQ), as well as SSL-MSE values calculated on the features obtained from the last layer of the SSL model. The performance of downstream tasks was evaluated on a subset of the SUPERB benchmark. Namely, we chose ASR, phone recognition (PR), automatic speaker verification (ASV), intent classification (IC), and slot-filling (SF) tasks to cover three major categories of tasks, i.e., content, speaker, and semantics. ASR performance was measured on the test-clean set and without the language model. To evaluate the noise robustness, we prepared a noisy version of the SUPERB test sets for each task by adding DNS noise to the original recordings at SNR values randomly sampled from 0 to 10 dB. Since DNS noise contains a wide variety of the noise types and we use different samples for training and testing, we can verify robustness to unseen noise although we adopted DNS noise in both training and evaluation.

### 4.2. Experimental results

#### 4.2.1. Performance comparison

Table 1 shows the performance of the SE model and the SSL pipeline with the SE frontend. Compared with the results gained from the clean test sets (a1), the addition of noise substantially degrades the performance of SSL pipeline as shown in (a2). Although WavLM BASE+ is trained robustly to DNS noise, the baseline SE frontend still improves the downstream performance in 4 out of 5 tasks (a3).

Column ‘SSL-MSE’ in the table indicates that by introducing the SE frontend trained with SSL-MSE (a4)-(a7), the output features of the SSL model for noisy speech become closer to those for clean speech compared with that of the baseline SE frontend (a3). Accordingly, the proposed systems considerably improved the performance of the downstream tasks compared with the baseline. The system (a4)-(a6) show that SSL-MSE training also slightly improved the SDR and PESQ compared with the baseline when the SNR weight of SSL-MSE training is set as 0.1. The system trained only with SSL-MSE loss (a7) shows the best performance in three out of five downstream tasks, while SDR and PESQ largely drop compared with system (a4)-(a6). This indicate that there is a tradeoff between SE performance and the compatibility with the SSL models, which we further discuss about in Section 4.2.2. As for the layer weight for the SSL-MSE, the latter-half setup shows su-

Table 1: The performance of each system evaluated by SE performance and downstream performance with **WavLM Base+**. For the downstream performance, row (a1) shows the result for official SUPERB eval sets and (a2)-(a7) for noisy eval sets. (a2) is the result of noisy speech without front-end processing. (a3) indicates performance with conventional SE model and (a4)-(a7) indicate that with SE trained with proposed SSL-MSE. Column ‘SSL-MSE’ shows the distance between enhanced or noisy speech and clean source measured on the SSL feature obtained from the last layer. SSL-MSE and the SE performance are evaluated on simulated noisy observation of Librispeech and DNS noise.

Test set	SE frontend	SSL MSE Loss	Layer weight $\tilde{w}$	SNR weight $\alpha$	Downstream					SSL MSE	Speech Enhancement		
					ASR	PR	ASV	IC	SF		SDR $\uparrow$	PESQ $\uparrow$	
					WER $\downarrow$	PER $\downarrow$	EER $\downarrow$	Acc $\uparrow$	F1 $\uparrow$				
(a1)	clean				5.6	4.4	4.5	98.8	90.3	-	-	-	
(a2)	noisy				17.1	16.0	10.9	67.4	82.4	0.0199	3.7	1.27	
(a3)	noisy	✓	No		14.6	9.9	8.6	84.8	82.4	0.0122	15.7	2.27	
(a4)	noisy	✓	Yes	last	0.1	13.1	8.8	8.0	87.6	83.5	0.0105	<b>15.8</b>	2.34
(a5)	noisy	✓	Yes	all	0.1	13.5	9.0	8.0	87.3	83.1	0.0106	<b>15.8</b>	<b>2.35</b>
(a6)	noisy	✓	Yes	latter-half	0.1	12.9	8.7	<b>7.9</b>	<b>87.8</b>	83.4	0.0103	<b>15.8</b>	<b>2.35</b>
(a7)	noisy	✓	Yes	last	0	<b>11.2</b>	<b>7.3</b>	8.4	87.7	<b>84.9</b>	<b>0.0093</b>	8.0	1.38

Table 2: The performance with SSL models other than that used for SSL-MSE training. ‘-’ in SE column means the performance without SE, ‘baseline’ is a condition with the SE trained on SNR loss, and ‘+SSL-MSE’ is the SE trained with SSL-MSE calculated on the last output of WavLM Base+ at SNR weight of 0.1.

SSL	SE	SSL MSE	Downstream			
			ASR	ASV	IC	
(b2)	w2v 2.0	-	0.1217	36.8	19.0	41.0
(b3)	w2v 2.0	baseline	0.0450	18.8	12.2	70.9
(b4)	w2v 2.0	+SSL-MSE	<b>0.0428</b>	<b>17.5</b>	<b>11.3</b>	<b>71.3</b>
(c2)	HuBERT	-	0.1207	31.7	15.5	57.8
(c3)	HuBERT	baseline	0.0696	18.5	10.6	81.9
(c4)	HuBERT	+SSL-MSE	<b>0.0628</b>	<b>16.9</b>	<b>9.7</b>	<b>82.6</b>
(d2)	WavLM L	-	0.0281	9.9	9.8	52.1
(d3)	WavLM L	baseline	0.0154	7.8	9.8	86.6
(d4)	WavLM L	+SSL-MSE	<b>0.0138</b>	<b>7.2</b>	<b>9.3</b>	<b>88.3</b>

Table 3: The performance for noise-robust downstream models. WavLM Base+ is used to obtain the result. (a2) in this figure corresponds to system (a2) in Table 1.

Noise robust downstream	SE	Downstream		
		ASR	ASV	IC
		WER $\downarrow$	EER $\downarrow$	Acc $\uparrow$
(a2)	-	17.1	10.9	67.4
(e2)	✓	13.6	9.2	79.4
(e3)	✓	12.5	7.9	85.8
(e4)	✓	<b>11.3</b>	<b>7.7</b>	<b>88.3</b>

perior performance on average (a6). Summarizing these results, the proposed SSL-MSE can successfully train the downstream task-agnostic SE model that improves the noise robustness of the SSL pipeline while maintaining the SE performance itself.

Table 2 shows the performance for SSL pipelines with SSL models other than WavLM Base+ which is used for SSL-MSE training. The results show that SSL-MSE improves the performance of the SSL pipeline even for the mismatched SSL models that are unseen during training, compared with the baseline SE training. Thus, it can be said that SSL-MSE training potentially generalizes over different types of SSL models.

Table 3 shows the performance of SSL pipelines when the downstream models are also trained to be robust to noise. As seen by comparing (a2) to (e2), the noise-robust training of the downstream model is effective for the noise-robust WavLM model. Row (e3) shows that the SE frontend is effective even with the noise-robust SSL model and downstream models. Furthermore, introducing the proposed SSL-MSE training (e4), achieves an additional performance improvement. This result indicates that SSL-MSE and the noise robust training of SSL model and downstream models are complementary and can be combined to obtain superior performance.




Figure 2: The tradeoff between enhancement and downstream ASR performance as a function of SNR loss weight. The results for WavLM Base+ and Wav2vec 2.0 are shown.

#### 4.2.2. Further discussion about the effect of SSL-MSE

To further understand the effect of the combination of SNR loss and SSL-MSE, we measured the SE and downstream ASR performance obtained with SE models trained with different SNR weights  $\alpha$ . Figure 2 shows that the downstream ASR performance with the WavLM Base+ model, shown by the red solid line, increases as SSL-MSE becomes even more dominant than SNR loss, while the performance of SE, shown by the blue solid line, decreases when SNR loss weight is less than 0.1. This indicates that SNR and SSL-MSE metrics requires different modes of enhancement and thus they are in a tradeoff relationship to some extent. The red dotted line plots the downstream ASR performance with Wav2vec 2.0 model. With a SSL model mismatched with that used for SSL-MSE calculation, the downstream performance decreases when SSL-MSE loss becomes significantly stronger than SNR loss. It seems that the SNR objective works as a regularization term and is required to maintain the generalizability of the SE model over different types of SSL models.

## 5. Conclusion

To construct an SE model that can effectively be combined with SSL models with less mismatch, we proposed SSL-MSE, which is a minimization criterion involving enhanced signals and clean sources in the SSL feature domain. The SE model trained with a multi-task loss including SSL-MSE is shown to improve the performance of downstream tasks in the combination with SSL models while retaining the SE performance itself. In addition to the downstream task-agnostic nature of the proposed method, it is shown that SSL-MSE can generalize over different types of SSL models.

## 6. References

- [1] D. Amodei, S. Ananthanarayanan, R. Anubhai, J. Bai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, Q. Cheng, G. Chen *et al.*, “Deep speech 2: End-to-end speech recognition in english and mandarin,” in *International conference on machine learning*. PMLR, 2016, pp. 173–182.
- [2] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath *et al.*, “Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups,” *IEEE Signal processing magazine*, vol. 29, no. 6, pp. 82–97, 2012.
- [3] A. Mohamed, H.-y. Lee, L. Borgholt, J. D. Havtorn, J. Edin, C. Igel, K. Kirchhoff, S.-W. Li, K. Livescu, L. Maaløe *et al.*, “Self-supervised speech representation learning: A review,” *IEEE Journal of Selected Topics in Signal Processing*, 2022.
- [4] X. Liu, F. Zhang, Z. Hou, L. Mian, Z. Wang, J. Zhang, and J. Tang, “Self-supervised learning: Generative or contrastive,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 1, pp. 857–876, 2021.
- [5] A. Jaiswal, A. R. Babu, M. Z. Zadeh, D. Banerjee, and F. Makeidon, “A survey on contrastive self-supervised learning,” *Technologies*, vol. 9, no. 1, p. 2, 2020.
- [6] Y.-A. Chung, C.-C. Wu, C.-H. Shen, H.-Y. Lee, and L.-S. Lee, “Audio Word2Vec: Unsupervised Learning of Audio Segment Representations Using Sequence-to-Sequence Autoencoder,” in *Proc. Interspeech*, 2016, pp. 765–769.
- [7] Y.-A. Chung and J. Glass, “Speech2Vec: A Sequence-to-Sequence Framework for Learning Word Embeddings from Speech,” in *Proc. Interspeech*, 2018, pp. 811–815.
- [8] S. Schneider, A. Baevski, R. Collobert, and M. Auli, “wav2vec: Unsupervised Pre-Training for Speech Recognition,” in *Proc. Interspeech*, 2019, pp. 3465–3469.
- [9] A. Baevski, S. Schneider, and M. Auli, “vq-wav2vec: Self-supervised learning of discrete speech representations,” in *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*, 2020.
- [10] Y.-A. Chung, Y. Zhang, W. Han, C.-C. Chiu, J. Qin, R. Pang, and Y. Wu, “W2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training,” in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 244–250.
- [11] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [12] 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 Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [13] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, “Wavlm: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [14] S.-w. 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 *et al.*, “SUPERB: Speech Processing Universal Performance Benchmark,” in *Proc. Interspeech*, 2021, pp. 1194–1198.
- [15] H.-S. Tsai *et al.*, “SUPERB-SG: Enhanced speech processing universal PERFORMANCE benchmark for semantic and generative capabilities,” in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*. ACL, May 2022, pp. 8479–8492.
- [16] Y. Masuyama, X. Chang, S. Cornell, S. Watanabe, and N. Ono, “End-to-end integration of speech recognition, dereverberation, beamforming, and self-supervised learning representation,” in *2022 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2023, pp. 260–265.
- [17] X. Chang, T. Maekaku, Y. Fujita, and S. Watanabe, “End-to-End Integration of Speech Recognition, Speech Enhancement, and Self-Supervised Learning Representation,” in *Proc. Interspeech*, 2022, pp. 3819–3823.
- [18] A. Pasad, J.-C. Chou, and K. Livescu, “Layer-wise analysis of a self-supervised speech representation model,” in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 914–921.
- [19] E. Dunbar, M. Bernard, N. Hamilakis, T. A. Nguyen, M. de Seyssel, P. Rozé, M. Rivière, E. Kharitonov, and E. Dupoux, “The Zero Resource Speech Challenge 2021: Spoken Language Modelling,” in *Proc. Interspeech*, 2021, pp. 1574–1578.
- [20] T.-A. Hsieh, C. Yu, S.-W. Fu, X. Lu, and Y. Tsao, “Improving Perceptual Quality by Phone-Fortified Perceptual Loss Using Wasserstein Distance for Speech Enhancement,” in *Proc. Interspeech 2021*, 2021, pp. 196–200.
- [21] J. Devlin, M. Chang, K. Lee, and K. Toutanova, “BERT: pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Volume 1*. Association for Computational Linguistics, 2019, pp. 4171–4186.
- [22] C. K. Reddy, H. Dubey, K. Koishida, A. Nair, V. Gopal, R. Cutler, S. Braun, H. Gamper, R. Aichner, and S. Srinivasan, “INTER-SPEECH 2021 Deep Noise Suppression Challenge,” in *Proc. Interspeech*, 2021, pp. 2796–2800.
- [23] K. Iwamoto, T. Ochiai, M. Delcroix, R. Ikeshita, H. Sato, S. Araki, and S. Katagiri, “How bad are artifacts?: Analyzing the impact of speech enhancement errors on ASR,” in *Proc. Interspeech*, 2022, pp. 5418–5422.
- [24] H. Sato, T. Ochiai, M. Delcroix, K. Kinoshita, N. Kamo, and T. Moriya, “Learning to enhance or not: Neural network-based switching of enhanced and observed signals for overlapping speech recognition,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 6287–6291.
- [25] H. Sato, T. Ochiai, M. Delcroix, K. Kinoshita, T. Moriya, and N. Kamo, “Should We Always Separate?: Switching Between Enhanced and Observed Signals for Overlapping Speech Recognition,” in *Proc. Interspeech*, 2021, pp. 1149–1153.
- [26] E. Vincent, R. Gribonval, and C. Févotte, “Performance measurement in blind audio source separation,” *IEEE transactions on audio, speech, and language processing*, vol. 14, no. 4, pp. 1462–1469, 2006.
- [27] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [28] Y. Luo and N. Mesgarani, “Conv-tasnet: Surpassing ideal time-frequency magnitude masking for speech separation,” *IEEE/ACM transactions on audio, speech, and language processing*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [29] M. Thakker, S. E. Eskimez, T. Yoshioka, and H. Wang, “Fast Real-time Personalized Speech Enhancement: End-to-End Enhancement Network (E3Net) and Knowledge Distillation,” in *Proc. Interspeech*, 2022, pp. 991–995.
- [30] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *International Conference on Learning Representations (ICLR)*, 2015.
- [31] <https://github.com/s3prl/s3prl> retrieved on 30/1/2023.
- [32] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, “Sdr-half-baked or well done?” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 626–630.