
Recycle-and-Distill: Universal Compression Strategy for Transformer-based Speech SSL Models with Attention Map Reusing and Masking Distillation

Kangwook Jang^{1*} Sungnyun Kim^{2*} Se-Young Yun² Hoirin Kim¹

¹School of Electrical Engineering, KAIST

²Graduate School of AI, KAIST

{dnrrkdwd12, ksn4397, yunseyoung, hoirkim}@kaist.ac.kr

Abstract

Transformer-based speech self-supervised learning (SSL) models, such as HuBERT, show surprising performance in various speech processing tasks. However, huge number of parameters in speech SSL models necessitate the compression to a more compact model for wider usage in academia or small companies. In this study, we suggest to reuse attention maps across the Transformer layers, so as to remove key and query parameters while retaining the number of layers. Furthermore, we propose a novel masking distillation strategy to improve the student model’s speech representation quality. We extend the distillation loss to utilize both masked and unmasked speech frames to fully leverage the teacher model’s high-quality representation. Our universal compression strategy yields the student model that achieves phoneme error rate (PER) of 7.72% and word error rate (WER) of 9.96% on SUPERB benchmark.

1 Introduction

Transformer-based speech SSL models [9, 19, 20] have been actively studied in speech processing field [21] as SSL arises as a successful representation learning approach in recent years [3, 5, 8, 22]. Especially for wav2vec 2.0 [2], HuBERT [16], and wavLM [7], all of which are inherited from BERT [10], show surprising performance in automatic speech recognition (ASR), comparable to supervised learning approaches [13, 17]. Since the versatility of speech SSL becomes also crucial, the above models have been further explored in various applications including automatic speaker verification (ASV) [30] or emotion recognition (ER) [25].

However, these models have huge number of parameters and are trained for very long time, which makes it hard for the resource-limited groups to train their own models. For instance, wav2vec 2.0 LARGE with 317M parameters should be pretrained for more than 290 days on a single V100 GPU [2] on LibriSpeech dataset [24]. This necessitates us to build a compressed model that allows much more parameter-efficient training and lower computational overhead.

Knowledge distillation (KD) [15] is a common model compression technique where a smaller student model is being trained by distilling the knowledge from a teacher model. Prior efforts in distilling large-scale speech SSL models have been made with reducing the number of Transformer layers or shrinking their width. DistilHuBERT [6] is distilled in a way of predicting multi-layer outputs of HuBERT, with most of the Transformer layers removed. FitHuBERT [18], instead of removing the layers, suggests cutting down the width of attention and feed-forward network (FFN) in each Transformer layer. LightHuBERT [29] creates a prunable supernet through distillation and conducts architecture search to make a small student.

*These authors contributed equally to this work.

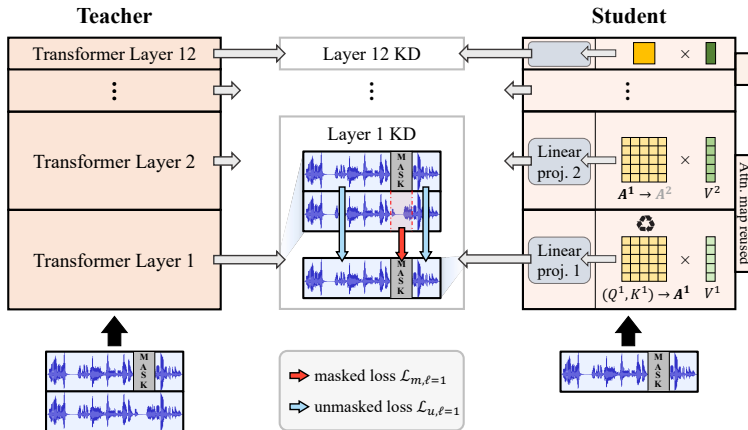


Figure 1: Our compression strategy involves reusing the attention map of the previous layer and extending the distillation process to masked (red arrow) and unmasked (blue arrow) representations. The input masked frames are identical for both teacher and student.

Despite the effectiveness of previous approaches in mitigating the performance drop by compression, they still face several issues. (1) Wide and shallow students [1, 6] still exhibit degradation on content-related downstream tasks. (2) Layer-to-layer (L2L) distillation is proved to be effective [1, 18], however, it is counter-intuitive in terms of compression since every layer’s parameters are required. (3) Pruning by architecture search [29] prepares an additional teacher-sized supernet using 32 GPUs, which is not end-to-end (E2E) and cannot be easily trained by resource-limited groups.

We suggest reusing attention maps across the student’s Transformer layers, which is inspired by previous works [4, 32] that claimed the similarity between attention maps. Attention map reusing enables us to remove key and query parameters in certain Transformer layers, making it unnecessary to retain all layer parameters for L2L distillation. Furthermore, we can reinvest the saved parameters to other parts of Transformer. We also propose a masking with L2L distillation for better speech representation quality of our student model. Masking speech frames is a widely used technique in speech SSL models [2, 16], trained by predicting the masked representation. This technique has been simply applied to distilling HuBERT [29], but not in the L2L manner. Our novel masking distillation scheme aims to fully leverage the teacher’s representation by extending the distillation loss to both masked and unmasked speech frames. We emphasize that our scheme is an E2E fashion and enhances the general quality of speech representation, especially in content- and semantics-related tasks.

Combining our two approaches described (Fig. 1), we reinvest the saved parameters from attention map reusing to FFN, and create our flagship model, **ARMHuBERT** (Atention map Reused Mask HuBERT). As evaluated on the SUPERB benchmark [31], ARMHuBERT achieves overall score [7] of 78.1, the state-of-the-art E2E distillation. It also reaches 7.72% PER in phoneme recognition (PR) task, and 9.96% WER in ASR task.

2 Preliminaries on Transformer-based Speech SSL Models

Recent dominant SSL models in speech field are wav2vec 2.0 [2], HuBERT [16], and wavLM [7], where these three model structures are identical except for detailed level. Specifically, they share 12 or 24 Transformer [28] layers and 7-layer 1D-CNN. Their pretraining schemes are based on masked prediction, estimating the codewords by output representations of the masked frames. Despite the superiority and scalability of speech SSL models, large number of parameters and their computational overhead make it difficult to train or utilize these models. We thus implement the model compression on HuBERT and wavLM, the two dominant SSL models in speech, to demonstrate the effectiveness of our compression strategy.

3 Methodology

3.1 Attention Map Reusing

Attention map reusing is a technique for substituting the present layer’s attention with the previous one, which has been covered in several domains [27, 32]. Prior works [4, 32] have pointed

out the similarity of the attention maps across heads and layers in pretrained Transformer models, such as BERT [10] and ViT [11]. We leverage this property by reusing the attention maps to compress the student model. Alternatively, we can reassign the amount of parameters saved by attention map reusing, without increasing the total number of parameters. Appx. A gives brief description and notations for multi-head self-attention (MHSA) module of Transformers.

Attention map reusing is to replace the attention map A_h with the previous layer’s one. For instance, if we reuse the k -th previous attention map on the current layer ℓ , the ReuseMHSA module is

$$\text{ReuseMHSA}(x) = [A_1^{\ell-k} V_1^\ell, \dots, A_H^{\ell-k} V_H^\ell] W_o^\ell. \quad (1)$$

Accordingly, computing K_h and Q_h can be omitted, reducing the number of multiplications and additions by $(2nd^2 + n^2d)$. Assuming $d/H = d_v = d_k$, the omitted computation accounts for half of the original computation for MHSA, which is $(4nd^2 + 2n^2d)$. Hence, less parameters and multiply-accumulates (MACs) are required as more ReuseMHSA modules are employed (see Appx. C.1).

3.2 Masking Distillation

Attention map reusing has reduced the number of parameters, however, it may affect the representation quality of the student model. To improve the student’s representation learning, we offer a novel masking distillation scheme that leverages the teacher’s representation knowledge in a more sophisticated way. Speech frame masking involves learning representation through masked prediction, where the model learns to represent masked frames accurately based on other unmasked frames. LightHuBERT [29], inspired by data2vec [3], has first applied the masking strategy to distilling HuBERT. In this approach, teacher model guides the representation of masked frames. Let $\mu(x)$ be a masked input, and f^t and f^s the teacher and student model. Then, the masked loss function becomes

$$\mathcal{L}(x) = \frac{1}{|M|} \sum_{i \in M} \|f_i^t(x) - f_i^s(\mu(x))\|_2 \quad (2)$$

where f_i is the i -th frame of the speech representation, and M is the set of the masked frames.

In addition to the masked part loss (eq. 2), we suggest to employ an unmasked loss since the teacher model can provide high-quality representation even on the unmasked frames. However, if the masking process removes essential frames, distilling the intact form of $f^t(x)$ can leak such essential knowledge that should have been removed. This induces biased predictions of the student, as it learns information that cannot be inferred from the masked input.

To prevent this, we make the teacher model receive the same masked input as the student does when distilling the unmasked part. Hence, the entire distillation loss becomes

$$\begin{aligned} \mathcal{L}(x) &= \sum_{\ell} \alpha_{\ell} [\mathcal{L}_{m,\ell}(x) + \mathcal{L}_{u,\ell}(x)] \\ &= \sum_{\ell} \frac{\alpha_{\ell}}{|M|} \sum_{i \in M} \|f_{i,\ell}^t(x) - f_{i,\ell}^s(\mu(x))\|_2 + \sum_{\ell} \frac{\alpha_{\ell}}{n - |M|} \sum_{i \notin M} \|f_{i,\ell}^t(\mu(x)) - f_{i,\ell}^s(\mu(x))\|_2 \end{aligned} \quad (3)$$

where α_{ℓ} is the layerwise coefficient. $\mathcal{L}_{m,\ell}$ and $\mathcal{L}_{u,\ell}$ represent masked loss and unmasked loss of the ℓ -th layer, respectively. In summary, our novel masking distillation strategy appropriately guides the student’s knowledge acquisition, by distilling not only the masked representation of unmasked data but also the unmasked representation of masked data (see Fig. 1). In Appx. C.2, we investigate the strength of our masking strategy compared to other types of losses.

4 Results

4.1 Model Description

To verify our masking distillation strategy, we first build a student model, MaskHuBERT, which employs masking distillation only. MaskHuBERT has the width of (attention, FFN) as (480, 640). Then, 2by6 reuse pattern is applied to MaskHuBERT, leading to 10.3% of parameter reduction. We extend this model to two options: ARMHuBERT and ARMHuBERT-S. ARMHuBERT is a reinvested version of MaskHuBERT, where the saved parameters from attention map reusing are reassigned to FFN, resulting in increased width of (480, 864). ARMHuBERT-S is a reduced version to match the

Table 1: Evaluation results on the SUPERB benchmark. For the details on this benchmark, refer to Appx. B.3. Metrics include parameter size in million, PER%, WER% (w/o language model), accuracy (Acc%), maximum term weighted value (MTWV), equal error rate (EER%), diarization error rate (DER%), F1 score (F1%), and concept error rate (CER%). ‘‘Overall’’ denotes the average scoring of all tasks proposed in [7]. LightHuBERT [29] works by two-stage training, where HuBERT-sized supernet needs to be trained first, thus not compared with E2E distillation models.

Models	Params Millions ↓	Overall Score ↑	Content				Speaker			Semantics			Paral.
			PR PER ↓	ASR WER ↓	KS Acc ↑	QbE MTWV ↑	SID Acc ↑	ASV EER ↓	SD DER ↓	IC Acc ↑	SF F1 ↑	CER ↓	ER Acc ↑
<i>Baselines</i>													
FBANK [31]	0	40.5	82.01	23.18	8.63	0.0058	8.5E-4	9.56	10.55	9.1	69.64	52.94	35.39
HuBERT BASE [16]	94.70	80.8	5.41	6.42	96.30	0.0736	81.42	5.11	5.88	98.34	88.53	25.20	64.92
wavLM BASE [7]	94.70	81.9	4.84	6.21	96.79	0.0870	84.51	4.69	4.55	98.63	89.38	22.86	65.94
LightHuBERT a_{Small} [29]	94.7 → 27.00	79.1	6.60	8.33	96.07	0.0764	69.70	5.42	5.85	98.23	87.58	26.90	64.12
<i>960h distillation – # params: 26.4M ~ 31.6M</i>													
FitW2V2 [18]	31.63	76.5	12.22	11.44	96.04	0.0475	64.71	6.65	6.44	93.38	86.65	29.40	62.35
3-L ONE-Pred [1]	30.58	76.8	13.34	12.23	96.69	0.0489	75.71	6.48	6.56	94.15	82.89	34.65	63.95
12-L HALF-L2L [1]	26.87	77.6	10.67	10.96	97.24	0.0604	69.52	6.13	6.81	96.97	86.11	30.93	63.24
MaskHuBERT (ours)	26.64	77.8	7.30	9.77	96.36	0.0664	62.83	5.38	6.79	97.05	87.31	27.10	62.37
ARMHuBERT (ours)	26.45	78.1	7.72	9.96	96.88	0.0635	65.03	5.68	7.10	97.07	87.59	26.06	62.86
<i>960h distillation – # params: 22.4M ~ 23.5M</i>													
DistilHuBERT [6]	23.49	75.9	16.27	13.37	95.98	0.0511	73.54	8.55	6.19	94.99	82.57	35.59	63.02
FitHuBERT [18]	22.49	74.5	13.32	12.09	96.27	0.0489	55.71	8.00	6.84	91.25	84.06	32.46	59.82
ARMHuBERT-S (ours)	22.39	77.5	8.63	10.82	96.82	0.0720	63.76	5.58	7.01	97.02	86.34	29.02	62.96
ARMwavLM-S (ours)	22.39	78.9	7.42	10.03	97.01	0.0741	71.29	5.99	7.11	97.76	87.41	26.97	64.54
<i>100h distillation</i>													
FitW2V2 [18]	22.49	73.1	16.50	14.77	94.68	0.0380	51.65	7.43	6.94	90.03	81.95	34.74	62.87
FitHuBERT [18]	22.49	74.5	14.05	12.66	96.23	0.0579	54.24	7.88	7.19	94.20	83.41	34.00	61.67
ARMHuBERT-S (ours)	22.39	76.8	9.17	11.83	96.01	0.0569	66.48	5.92	6.23	95.97	83.89	33.29	63.29
ARMwavLM-S (ours)	22.39	77.0	8.33	11.37	96.30	0.0579	65.40	6.38	7.41	96.76	84.89	31.95	63.41

parameters with previous works, having the width of (432, 816). To establish the universality of our strategy, we introduce ARMwavLM-S that is structurally identical to ARMHuBERT-S, with the only change in teacher from HuBERT [16] to wavLM [7]. Refer to Appx. B for implementation details.

4.2 SUPERB Benchmark Results

In Table 1, we evaluate our student models on the SUPERB benchmark [31]. We follow the default fine-tuning recipes, including a learning rate scheduler, with the learning rate scaled to $10\times$ in SID task. MaskHuBERT outperforms 12-L HALF-L2L, the previous state-of-the-art E2E distillation method, with less parameters used. Our observation indicates that incorporating our masking strategy into the L2L distillation [1, 18] results in enhancing the student’s representation quality. Especially, MaskHuBERT highly improves the performances in content- and semantics-related tasks.

ARMHuBERT achieves a better overall score of 78.1 with less parameters than MaskHuBERT. Despite the removal of certain attention parameters, increasing the FFN width contributes to better quality of speech representation, achieving 7.72% PER and 9.96% WER. We find out that ARMHuBERT shows promising improvements when compared to MaskHuBERT in SF and SID tasks, exhibiting a similar level of performance in other tasks. In the end, the number of parameters and MACs in ARMHuBERT have decreased to 28% and 30% of the teacher model, HuBERT BASE [16].

In a smaller parameter group, ARMHuBERT-S, the parameter-reduced version, outperforms DistilHuBERT and FitHuBERT by a large margin. Specifically, ARMHuBERT-S also shows the outstanding results in content- and semantics-related tasks, which means the consistency of the representations produced by MaskHuBERT and ARMHuBERT-S. In addition, the result that ARMwavLM-S surpasses ARMHuBERT-S implies the universality of our strategy: without any modifications of the student model structure, replacing with a superior teacher model creates a better student. The results of the LibriSpeech [24] 100h distillation are also consistent with the formerly demonstrated results.

5 Conclusion and Future Work

In summary, we have proposed the universal compression strategy which involves attention map reusing and novel masking distillation. Our parameter-reinvested model, ARMHuBERT, achieves great performance in content- and semantics-related tasks. Our strategy can be applied to any Transformer-based speech SSL models, and contributes to enhancing the general quality of speech representation. Future work can focus on further improving our model on speaker-related tasks.

Acknowledgments and Disclosure of Funding

The study was supported by Korea Health Technology R&D Project through the Korea Health Industry Development Institute funded by the Ministry of Health and Welfare, Republic of Korea (HR18C0016).

References

- [1] T. Ashihara, T. Moriya, K. Matsuura, and T. Tanaka. Deep versus wide: An analysis of student architectures for task-agnostic knowledge distillation of self-supervised speech models. In *Proc. Interspeech*, 2022. 2, 4
- [2] 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*, 33: 12449–12460, 2020. 1, 2
- [3] 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 *Proc. ICML*, pages 1298–1312. PMLR, 2022. 1, 3
- [4] S. Bhojanapalli, A. Chakrabarti, A. Veit, M. Lukasik, H. Jain, F. Liu, Y.-W. Chang, and S. Kumar. Leveraging redundancy in attention with reuse transformers. *arXiv preprint arXiv:2110.06821*, 2021. 2
- [5] M. Caron, P. Bojanowski, A. Joulin, and M. Douze. Deep clustering for unsupervised learning of visual features. In *Proc. ECCV*, pages 132–149, 2018. 1
- [6] H.-J. Chang, S.-w. Yang, and H.-y. Lee. Distilhubert: Speech representation learning by layer-wise distillation of hidden-unit bert. In *Proc. ICASSP*, pages 7087–7091. IEEE, 2022. 1, 2, 4
- [7] 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*, 16(6):1505–1518, 2022. 1, 2, 4, 7
- [8] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. In *Proc. ICML*, pages 1597–1607. PMLR, 2020. 1
- [9] P.-H. Chi, P.-H. Chung, T.-H. Wu, C.-C. Hsieh, Y.-H. Chen, S.-W. Li, and H.-y. Lee. Audio albert: A lite bert for self-supervised learning of audio representation. In *Proc. SLT*, pages 344–350. IEEE, 2021. 1
- [10] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proc. NAACL*, pages 4171–4186, 2019. 1, 3
- [11] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proc. ICLR*, 2021. 3
- [12] T.-h. Feng, A. Dong, C.-F. Yeh, S.-w. Yang, T.-Q. Lin, J. Shi, K.-W. Chang, Z. Huang, H. Wu, X. Chang, et al. Superb@ slt 2022: Challenge on generalization and efficiency of self-supervised speech representation learning. In *Proc. SLT*, pages 1096–1103. IEEE, 2023. 8
- [13] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, et al. Conformer: Convolution-augmented transformer for speech recognition. In *Proc. Interspeech*, 2020. 1
- [14] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick. Masked autoencoders are scalable vision learners. In *Proc. CVPR*, pages 16000–16009. IEEE, 2022. 8
- [15] G. Hinton, O. Vinyals, J. Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015. 1

- [16] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, 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*, 29:3451–3460, 2021. 1, 2, 4, 7, 8
- [17] S. Kim, A. Gholami, A. E. Shaw, N. Lee, K. Mangalam, J. Malik, M. W. Mahoney, and K. Keutzer. Squeezeformer: An efficient transformer for automatic speech recognition. *Advances in Neural Information Processing Systems*, 2022. 1
- [18] Y. Lee, K. Jang, J. Goo, Y. Jung, and H. Kim. Fithubert: Going thinner and deeper for knowledge distillation of speech self-supervised learning. In *Proc. Interspeech*, 2022. 1, 2, 4, 7
- [19] 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 *Proc. ICASSP*, pages 6419–6423. IEEE, 2020. 1
- [20] A. T. Liu, S.-W. Li, and H.-y. Lee. Tera: Self-supervised learning of transformer encoder representation for speech. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:2351–2366, 2021. 1
- [21] 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*, 3(12):100616, 2022. 1
- [22] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013. 1
- [23] A. v. d. Oord, Y. Li, and O. Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 7
- [24] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur. Librispeech: an asr corpus based on public domain audio books. In *Proc. ICASSP*, pages 5206–5210. IEEE, 2015. 1, 4, 7, 8
- [25] L. Pepino, P. Riera, and L. Ferrer. Emotion Recognition from Speech Using wav2vec 2.0 Embeddings. In *Proc. Interspeech*, pages 3400–3404, 2021. 1
- [26] S. Schneider, A. Baevski, R. Collobert, and M. Auli. wav2vec: Unsupervised pre-training for speech recognition. In *Proc. Interspeech*, 2019. 7
- [27] K. Shim, J. Choi, and W. Sung. Understanding the role of self attention for efficient speech recognition. In *Proc. ICLR*, 2022. 2
- [28] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2, 7
- [29] R. Wang, Q. Bai, J. Ao, L. Zhou, Z. Xiong, Z. Wei, Y. Zhang, T. Ko, and H. Li. Lighthubert: Lightweight and configurable speech representation learning with once-for-all hidden-unit bert. In *Proc. Interspeech*, 2022. 1, 2, 3, 4, 8
- [30] Y. Wang, A. Boumadane, and A. Heba. A fine-tuned wav2vec 2.0/hubert benchmark for speech emotion recognition, speaker verification and spoken language understanding. *arXiv preprint arXiv:2111.02735*, 2021. 1
- [31] S. wen Yang, P.-H. Chi, Y.-S. Chuang, C.-I. J. Lai, K. Lakhotia, 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 *Proc. Interspeech 2021*, pages 1194–1198, 2021. 2, 4, 7
- [32] T. Xiao, Y. Li, J. Zhu, Z. Yu, and T. Liu. Sharing attention weights for fast transformer. In *Proc. IJCAI*, pages 5292–5298, 2019. 2

Appendix

A Multi-head Self-attention of Transformers

In Transformer’s multi-head self-attention (MHSA) module [28], the input $x \in \mathbb{R}^{n \times d}$ with the sequence length n is transformed to H independent queries, keys, and values by transformation matrices $W_{h,k}, W_{h,q} \in \mathbb{R}^{d \times d_k}$, and $W_{h,v} \in \mathbb{R}^{d \times d_v}$, respectively, for each head h . Here, d_k, d_v , and d are the width of the keys, values, and model, respectively.

$$\begin{aligned} K_h &= W_{h,k}x, & K_h &\in \mathbb{R}^{n \times d_k}, \\ Q_h &= W_{h,q}x, & Q_h &\in \mathbb{R}^{n \times d_k}, \\ V_h &= W_{h,v}x, & V_h &\in \mathbb{R}^{n \times d_v} \end{aligned} \tag{4}$$

Then, key and query are multiplied along the width axis to obtain a scaled dot-product attention map, $A_h \in \mathbb{R}^{n \times n}$. Linear combinations of the attention map and value for each head are concatenated, and then projected to the original width.

$$A_h = \text{softmax}(Q_h K_h^\top / \sqrt{d_k}), \tag{5}$$

$$\text{MHSA}(x) = [A_1 V_1, \dots, A_H V_H] W_o, \quad W_o \in \mathbb{R}^{H d_v \times d} \tag{6}$$

B Implementation Details

B.1 Training

We distilled the two dominant Transformer-based speech SSL models, HuBERT BASE [16] and wavLM BASE [7], that are pretrained on LibriSpeech 960 hours dataset [24]. Our student model consists of 12 layers of Transformers as the teachers, while the detailed design mostly follows FitHuBERT [18]: width of attention and FFN reduced and linear projections adopted at each layer. The layerwise coefficients α_ℓ are set to 0.1 except for the last layer, where it is set to 1. Unless specified, the LibriSpeech [24] dataset is distilled for 200 epochs with effective batch size of 72 including gradient accumulation.

B.2 Reuse pattern

We employ an alternating reuse pattern for the attention maps, whereby the attention map of an even-numbered Transformer layer is repeated by that of the previous odd-numbered layer. We denote this pattern as *2by6*, our default setting. We examine other reuse patterns in Appx. C.1 in terms of performance, number of parameters, and MACs.

B.3 SUPERB Benchmark

The beginning of speech SSL models focused on content-related downstream tasks such as ASR or PR [23, 26], however, their versatility to other tasks has been recognized as crucial recently [7]. In this context, SUPERB benchmark [31] has been proposed to evaluate the generalizability of speech SSL models, covering the aspects of content, speaker, semantics, and paralinguistics. We evaluate our representation against the SUPERB benchmark to verify the generalizability of our student model. The SUPERB downstream tasks include PR, ASR, keyword spotting (KS), query-by-example spoken term detection (QbE), speaker identification (SID), ASV, speaker diarization (SD), intent classification (IC), slot filling (SF), and ER.

C Discussions on Experiments

In this section, we explore which layer’s attention map should be reused in other layers and how to implement the masking distillation. Unless specified, we have conducted the distillations on 100-hour of LibriSpeech [24] and evaluated on the ASR, ASV, and SF tasks of the SUPERB benchmark [31].

C.1 Where to Reuse

Table 2 summarizes the performance depending on various attention map reusing patterns, and in general, the *2by6* pattern performs the best. Other reuse patterns have reduced the Transformer’s representation capacity due to overly frequent reusing. Assigning more parameters to FFN (*-up*) still has limit in terms of the performance gain. Comparing to no reuse pattern applied, the performance decrease of *2by6* is small, but it takes advantages in 9.13% and 8.16% reduction of parameters and MACs, respectively. We note that the number of MACs in a single reuse MHSA module (eq. 1) is reduced by half, from 13.2G to 6.6G.

Table 2: Performance comparisons of various reusing patterns. Parameter size (M) and MACs (G) are additionally measured [12]. The width of (attention, FFN) for each model is (432, 816), while “-up” suffix denotes more parameters assigned to FFN to match with *2by6*. Masking is not applied here.

pattern	reused layers	params	MACs	WER ↓	EER ↓	F1 ↑	CER ↓
<i>6by2</i>	{1,7}	20.90	423	13.52	6.30	83.69	34.92
<i>3by4</i>	{1,4,7,10}	21.65	437	12.37	5.67	83.29	33.60
<i>6by2-up</i>	{1,7}	22.39	440	13.18	5.89	83.07	34.79
<i>3by4-up</i>	{1,4,7,10}	22.39	445	12.39	6.06	83.79	33.49
<i>2by6</i>	{1,3,5,7,9,11}	22.39	450	12.18	5.95	84.91	32.29
None	-	24.64	490	11.94	5.87	84.78	31.38

C.2 How to Mask

Masking strategy Table 3 shows the efficacy of our masking strategy. We first eliminated the loss function on the unmasked frames ($\mathcal{L}_{u,\ell}$), making it equivalent to the L2L version of the LightHuBERT [29] distillation loss. This approach severely damaged performances, particularly in the ASR and ASV tasks. Next, we modified the unmasked loss function to distill from the unmasked input, *i.e.*, only $f^t(x)$ being distilled to the student. This also led to degraded performance in most tasks, revealing that our unmasked loss with masked input properly guides the knowledge acquisition without imposing biased predictions.

Table 3: Ablation study on our masking strategy.

methods	WER ↓	EER ↓	F1 ↑	CER ↓
MaskHuBERT-100h	11.56	5.87	84.31	32.28
[-] distil. unmasked part	13.23	7.96	82.78	33.53
[-] distil. from masked input	11.65	6.09	84.29	31.41

Masking ratio High value of masking ratio can lead to a student model producing good representation, as it has less information to infer with [14, 16]. However, it can also make the learning process more difficult. In Table 4, we examine the optimal masking ratios for each training set. For LibriSpeech 960h [24], both ratios of 0.4 and 0.8 produce excellent results. On the other hand, for the 100h dataset, ratio of 0.4 produces the best results overall. This implies that lower masking ratio is preferred in low-resource distillation setting. Accordingly, in our main experiments, we have used the ratios of 0.8 and 0.4 for the 960h and 100h distillation, respectively.

Table 4: Performance comparisons with different masking ratios. “sch” indicates linear scheduling of the ratio as 0.4 to 0.8.

models	ratio	WER ↓	EER ↓	F1 ↑	CER ↓
MaskHuBERT-960h	0.4	9.75	5.58	86.94	26.79
MaskHuBERT-960h	0.8	9.77	5.38	87.31	27.10
MaskHuBERT-100h	0.4	11.56	5.87	84.31	32.28
MaskHuBERT-100h	0.6	11.99	6.18	83.42	33.31
MaskHuBERT-100h	0.8	12.74	6.56	83.68	33.82
MaskHuBERT-100h	sch	12.07	6.29	83.84	33.50