

COCKTAIL HUBERT: GENERALIZED SELF-SUPERVISED PRE-TRAINING FOR MIXTURE AND SINGLE-SOURCE SPEECH

Maryam Fazel-Zarandi and Wei-Ning Hsu

Meta FAIR

{maryamfazel, wnhsu}@meta.com

ABSTRACT

Self-supervised learning leverages unlabeled data effectively, improving label efficiency and generalization to domains without labeled data. While recent work has studied generalization to more acoustic/linguistic domains, languages, and modalities, these investigations are limited to single-source speech with one primary speaker in the recording. This paper presents Cocktail HuBERT, a self-supervised learning framework that generalizes to mixture speech using a masked pseudo source separation objective. This objective encourages the model to identify the number of sources, separate and understand the context, and infer the content of masked regions represented as discovered units. Cocktail HuBERT outperforms state-of-the-art results with 69% lower WER on multi-speaker ASR, 31% lower DER on diarization, and is competitive on single- and multi-speaker tasks from SUPERB. Codes and mores are available at <https://github.com/facebookresearch/fairseq/tree/main/examples/hubert>.

Index Terms— Self-supervised pre-training, diarization, multi-speaker ASR, source separation, cocktail party, mixture speech

1. INTRODUCTION

Self-supervised learning (SSL) has greatly advanced speech processing over the past few years [1, 2, 3, 4, 5]. Supervised fine-tuning from a pre-trained model enjoys better label efficiency, achieving performance on par with supervised models using hundredths fewer labeled data [6]. The representations learned with pre-trained models are more universal: in contrast to those from supervised learning [7, 8, 9], self-supervised representations benefit a wider range of tasks [10, 11]. Self-supervised representations also enable many novel applications, such as unsupervised speech recognition and synthesis [12, 13, 14], disentangled speech codec [15], text-free spoken language models and prompting [16, 17].

A key advantage of self-supervised pre-training is that it uses unlabeled data instead of labeled data, such that a model can be pre-trained on data covering more domains [18, 19, 20, 21]. Consequently, the fine-tuned model is more robust to domain shift, suffering milder degradation when evaluated on domains unseen during fine-tuning [18]. Generalizing this idea, recent work also extends self-supervised pre-training to multi-modal speech [22, 23] and demonstrates a multi-modal speech recognition system can be built with only labeled unimodal data. However, up until now, speech pre-training has been designed for single-source speech, which contains one primary speaker in each sample whereas other sources are assumed noise [5], leaving mixture speech unattended.

Mixture speech, where multiple speakers may speak at the same time, occurs frequently in conversational scenarios. These scenarios impose greater challenges to applications common to single-source

speech (e.g., recognizing the “target” speech from a mixture), and also generate applications specific to mixture speech, including speech diarization, source separation, multi-speaker speech recognition (transcribe everything) and more. Pre-trained models designed for single-source speech are likely to be sub-optimal for these tasks.

In an effort to broaden the applicability of pre-trained models to wider speech varieties, this paper presents Cocktail HuBERT (C-HuBERT), a self-supervised framework that pre-trains on both single-source and mixture speech with a unified objective. The pre-training objective can be summed as *masked pseudo source separation*, which predicts automatically discovered units of randomly masked spans for each source in the mixture given unmasked speech context. The speech mixture contains one or more sources, created by artificially mixing single source samples. To excel at this task, the model is required to perform three tasks jointly: source separation, acoustic modeling, and language modeling. Evaluation on multi-speaker automatic speech recognition (MS-ASR) and speech diarization (SD) verifies that the proposed objective is particularly effective for downstream tasks concerning mixture speech, outperforming state-of-the-art results by large margins (7.8% vs 24.9% WER on Libri2Mix MS-ASR, and 3.86% vs 5.62% DER on Libri(2+3)Mix SD). Evaluation on SUPERB [10] also shows strong performance on additional mixture tasks and slight-to-none degradation on single-source tasks.

2. BACKGROUND

This paper is built upon Hidden Unit BERT (HuBERT) [3], one of the state-of-the-art speech pre-training frameworks. The pre-training objective of HuBERT is masked cluster prediction. Similar to BERT, it masks part of the speech input, and predicts given the context (unmasked part) some label derived from the masked input. While the label used in BERT is the input token itself, HuBERT proposes to obtain discrete labels via clustering audio features and refine the label iteratively. Concretely, let y be waveform, $x_t = f_t(y)$ be HuBERT local feature extractor f (CNN) output at time t , $c_t^l = g_t^l(x)$ be contextualized feature extractor g (L -layer Transformer) output at time t layer l , and z_t be the target unit at time t . HuBERT (f and g) is pre-trained by predicting z_t from $g_t^L(\text{MASK}(x))$ for time steps t that are masked, where $\text{MASK}(\cdot)$ is an operation that randomly samples spans of 10 frames and replaces the features in those spans with a learned masked embedding [MSK] following wav2vec2.0 [6]. In the first iteration, z_t are obtained by clustering MFCC features. In the subsequent iterations, the latest iteration HuBERT representations c_t^L are used for clustering, which produces higher quality cluster assignments than those from raw or earlier iteration HuBERT features.

Intuitively, HuBERT pre-training solves acoustic and language modeling task jointly, where the model needs to understand the con-

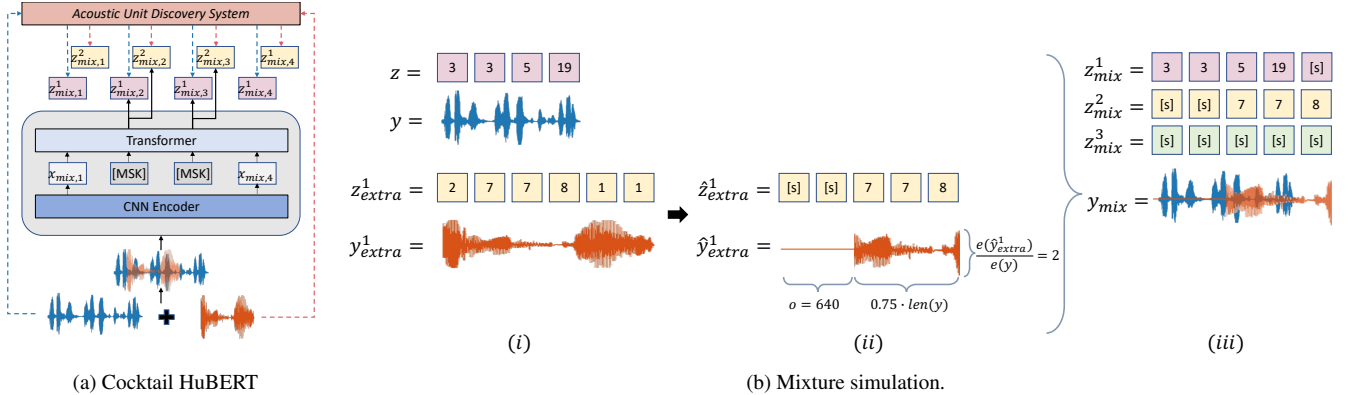


Fig. 1: (a) C-HuBERT ($K = 2$) predicts hidden units of the masked frames for each source in the input audio mix generated by k-means clustering. (b) Mixture simulation with $K = 3$, $n = 1$, $(r_l, r_e, o) = (0.75, 2.0, 640)$. Step i: sample the number of extra sources n and then sample additional sources $z_{extra}^{1:n}$. Step ii: chunk, scale, and shift according to sampled (r_l, r_e, o) . $e(y)$ denotes the energy of y . Step iii: mix audio and pad target units with [SIL] for silent frames (last frame of z_{mix}^1 and first two frames of z_{mix}^2) and silent streams (z_{mix}^3).

tent from observed regions (acoustic model) and then infer the label for the masked frames (language model).

3. COCKTAIL HUBERT

The ‘‘cocktail party problem’’ describes the setup where sounds from different sources are mixed prior to being perceived, such that estimation of individual sources given the mixture is required to perform downstream tasks. Human brains have impressive ability to focus on a particular stimulus leveraging structural properties of single sources. Researchers have also attempted to reproduce this capability and develop applications like source separation [24] and multi-speaker speech recognition [25].

In this paper, we present Cocktail HuBERT, a self-supervised learning objective that trains the model to tackle the cocktail party problem, as depicted in Figure 1a. Conceptually, the model takes as input a masked mixture which contains one or more speech sources with spans of frames randomly masked, similar to HuBERT. The model processes the masked mixture and predicts the discrete labels for each source for only the masked frames. The discrete labels can be viewed as automatically discovered frame-level phonetic transcripts. Thus, the pre-training task is analogous to pseudo source separation that predicts a proxy frame representation for each source.

3.1. Model and Training Objective

Cocktail HuBERT adopts the same local and contextualized feature encoder (f and g) as HuBERT, but instead of having only one projection head to predict the posterior over units for a single source, Cocktail HuBERT has K project heads to predict units for at most K sources. Let y_{mix} be a mixture containing up to K sources and z_{mix}^i for $i \in [K]$ be the target unit sequence for source i (some corresponds to silent sources). The model outputs K streams of predictions, where $p_t^j(\cdot | g(\text{MASK}(f(y_{mix}))))$ denotes the j -th stream prediction at time step t . The loss of the j -th stream prediction with respect to the i -th source unit sequence z_{mix}^i is:

$$L_m^{j,i} = \sum_{t \in M} \log p_t^j(z_{mix,t}^i | g(\text{MASK}(f(y_{mix})))), \quad (1)$$

where M denotes the masked time steps. Since the order of the model predictions j do not necessarily align with the order of the sources i , permutation invariant training (PIT) loss [26] is deployed

which finds the alignment with the minimum loss. Let \mathcal{P} be all permutations of $[K]$, the masked pseudo source separation objective is

$$\frac{1}{K} \arg \min_{\pi \in \mathcal{P}} \sum_{j=1}^K L_m^{j,\pi(j)}. \quad (2)$$

3.2. Mixture Simulation

For training our models, overlapped speech for a maximum of K speakers is simulated as follows (See Fig 1b). A batch of B utterances where $B \geq K$ is sampled from the dataset. For each utterance y and its units z , the number of additional sources $n \in \{0, \dots, K-1\}$ is sampled with $P(n=0) = 1 - p_{mix}$ and $P(n=k) = p_{mix}/(K-1)$ for $\forall k \neq 0$. An additional source is either an utterance from the batch, or a non-speech noise sampled from a noise dataset. The probability of selecting noise is p_{noise} . For each additional source y_{extra}^k , a tuple of length ratio, energy ratio, and offset (r_l, r_e, o) are sampled from some uniform distributions, which are used to chunk, scale, and shift y_{extra}^k with respect to y . Let \hat{y}_{extra}^k be the resulting source and \hat{z}_{extra}^k be the units chunked correspondingly if y_{extra}^k is not noise. The resulting mixture $y_{mix} = y + \sum_{k=1}^n \hat{y}_{extra}^k$. Note that each source is right-padded to the maximum length among y and $y_{extra}^k \forall k$ with silence. A special token [SIL] is used for frames corresponding to padded silence (including the offset at the beginning). The first $n+1$ unit sequences correspond to the [SIL]-padded z and \hat{z}_{extra}^k for non-noise k . The remaining $K - (n+1)$ sequences and those corresponding to noise samples are set to [SIL] sequences.

4. RELATED WORK

Cocktail HuBERT is most related to HuBERT [3] and WavLM [5], which are self-supervised speech pre-training frameworks based on masked cluster prediction. Similar to Cocktail HuBERT, WavLM also stochastically mixes single source speech with noise and/or other single source speech. However, WavLM is effectively HuBERT with data augmentation, which pre-trains the model using the same objective as HuBERT where only the units of the ‘‘primary’’ speech are predicted: the added noise and speech are both treated as noise and should be ignored by the model. Consequently, for the model to differentiate primary speech from interfering speech, WavLM deploys a more restrictive mixing strategy, where the interfering audio is at most half as long as the primary speech.

On the other hand, Cocktail HuBERT and [24] also share similar training objectives. Inspired by the capability of converting units back to speech [15] and the connection between single/multi-speaker speech recognition to speech enhancement/separation, [24] proposes an alternative approach to speech enhancement and source separation by predicting units instead of spectral masks or waveforms. Comparing Cocktail HuBERT with [24], the former are designed for pre-training, which masks partial input and predicts units only for masked frames, requiring the model to perform language modeling. It is also evaluated on a wide range of downstream tasks including both mixture and single-source speech processing. In contrast, the latter is designed for particular downstream tasks (enhancement and separation) that predict units for all frames without masking the input. It is unclear how the resulting model performs on other tasks.

5. EXPERIMENTAL SETUP

For unsupervised pre-training, we use 960 hours of LibriSpeech audio [27] for the BASE model and 60k hours of Libri-light audio [28] for the LARGE model. We extract features from the 9-th transformer layer of the HuBERT BASE model for K-means with 500 clusters and use those labels to train the Cocktail HuBERT models. This ensures that we have high quality labels. We apply data mixing to 20%, 60%, and 100% of the data, where the mixing noise probability is set to 10%. The Cocktail BASE and LARGE models are trained on 32 and 128 GPUs, respectively, for 400k and 800k steps. The batch sizes are at most 87.5 and 56.25 seconds of audio per GPU. Adam [29] optimizer is used with $\beta = (0.9, 0.98)$, and learning rate ramps up linearly from 0 to peak learning for the first 32k training steps, and then decays back to zero. The peak learning rates are $5e-4/1e-4$ for Cocktail BASE/LARGE models. K/p_{mix} are set to 5/1.0 for BASE and 3/1.0 for LARGE unless otherwise specified.

We evaluate our pre-trained models specifically on multi-speaker speech recognition and diarization tasks, as they involve multiple speakers. We use the LibriMix data [30] which contains multi-speaker overlapped speech simulated by mixing utterances from the LibriSpeech corpus. We focus on the two- and three-speaker scenarios, and mixes with the “max mode” where the shortest utterance is padded to the longest one. For MS-ASR, we fine-tune our models on multi-speaker labeled data using the connectionist temporal classification (CTC) [31] loss for each (output stream, label stream) pair and compute the PIT-CTC loss, to fine-tune the whole model except for the local feature extractor. The projection layers are removed and replaced with a randomly initialized softmax layer for each stream. The CTC target vocabulary includes 26 English characters, a space token, an apostrophe, and a special CTC blank symbol. We fine-tune each model on 8 GPUs on the `train-100-mix-clean` subset of Libri2Mix for the 2-speaker scenario. The batch sizes per GPU are at most 200/80 seconds of audio for BASE/LARGE models. We sweep over peak learning rate ($[1e-5, 1e-4]$) for each model size using the PIT word error rate (WER) on the `dev-mix-clean` subset as criterion for model selection. All other hyperparameters are based on [3], except that we set *freeze-step* to zero. We use beam search decoding with a 4-gram language model and a beam size of 500.

For SD, we use a similar setup to SUPERB [10], where we freeze the pre-trained model and weight-sum the representations from different g layers with learnable weights as the input to the diarization model. The diarization model uses a single layer 512-unit LSTM and is trained with the PIT loss. We train each model on 1 GPU on the `train-100-mix-both` subset of Libri2Mix and Libri3Mix, and a mixture of both datasets. We use a batch size of 8,

train for 30k steps, and sweep over learning rate ($[1e-2, 1e-4]$) for each model using accuracy on the `dev-mix-both` subset of each dataset as criterion for model selection. For the evaluation metric, we use the diarization error rate (DER) [32].

We also compare our models to other strong pre-trained models on a subset of SUPERB tasks [10, 11], including Phoneme Recognition (PR), Automatic Speech Recognition (ASR), Keyword Spotting (KS), Query by Example Spoken Term Detection (QbE), Intent Classification (IC), Slot Filling (SF), Emotion Recognition (ER), Speech Enhancement (SE), and Speech Separation (SS) following the protocol. Overall score is computed following [5].

6. RESULTS

6.1. Multi-speaker and Single-speaker ASR

We first evaluate Cocktail HuBERT models on multi-speaker ASR and compare them to three state-of-the-art supervised baselines: (1) the end-to-end ASR model trained with PIT-CTC [34], (2) the end-to-end ASR model trained with the extended Graph-based temporal classification [25] loss (GTC-e), and (3) the Conditional-Conformer-CTC model [33] that generates CTC predictions conditioning on past CTC predictions for other streams. To understand how pre-training objectives affect the multi-speaker ASR performance, we also fine-tune HuBERT BASE and LARGE, which have the identical model architecture as C-HuBERT models, as our self-supervised baselines.

Results are reported in Table 1. First, we observe that both Cocktail HuBERT BASE and LARGE significantly outperform the baselines, with LARGE reducing the WER by 69% relative (24.9% \rightarrow 7.8%). More importantly, there is a considerable gap between the HuBERT and Cocktail HuBERT performance (37.2% vs 13.7% for BASE, and 35.2% vs 7.8% for LARGE), validating that the proposed masked pseudo source separation objective brings significant gain to the multi-speaker downstream task.

We further investigate the performance of our MS-ASR models on single speaker input by comparing with models fine-tuned for single-speaker ASR (Table 2). Overall, we see degradation in performance across all settings when using MS-ASR models to transcribe single-speaker utterances. However, compared to HuBERT, C-HuBERT models are more robust to variation in the number of speakers: the WER of LARGE increases by 1.2% and 4.1% on test-

Table 1: WERs of multi-speaker ASR models trained on Libri2Mix and tested on `dev-mix-clean` (`dev-mix`) and `test-mix-clean` (`test-mix`). C-HuBERT BASE $p_{mix} = 0.6$.

Model	LM	Libri2Mix	
		dev-mix	test-mix
PIT-CTC [25]	Transformer	24.0	26.3
GTC-e [25]	Transformer	32.7	33.7
Cond-Conformer-CTC [33]	greedy	24.5	24.9
HuBERT BASE	4-gram	35.8	37.2
HuBERT LARGE	4-gram	34.4	35.2
C-HuBERT BASE	4-gram	12.7	13.7
C-HuBERT LARGE	4-gram	6.6	7.8

Table 2: WERs of multi-speaker and single-speaker ASR tested on single-speaker LibriSpeech sets: $\{\text{dev}, \text{test}\} - \{\text{clean}, \text{other}\}$ ($\{\text{d}, \text{t}\} - \{\text{c}, \text{o}\}$). C-HuBERT BASE $p_{mix} = 0.6$. 4-gram LMs are used

Model	finetune on LS-10h				finetune Libri2Mix			
	d-c	d-o	t-c	t-o	d-c	d-o	t-c	t-o
HuBERT B	4.1	9.4	4.5	9.7	9.1	20.7	10.1	23.6
C-HuBERT B	4.4	10.3	4.9	10.9	4.9	13.0	6.3	13.7
HuBERT L	2.5	5.2	3.0	5.6	5.8	14.3	8.1	14.7
C-HuBERT L	2.8	5.9	3.3	6.4	3.7	9.6	4.5	10.5

Table 3: Universal Speech Representation Evaluation on SUPERB. Results from all other models are from [5]. BASE/LARGE models are pre-trained on LS 960hr/Libri-light 60k hr, except for WavLM LARGE† which was trained on Mix 94k hr

Method	#Params	Content				Semantics			ParaL	Generation			Overall
		PR	ASR	KS	QbE	IC	SF		ER	SE		SS	Score ↑
		PER ↓	WER ↓	Acc ↑	MTWV ↑	Acc ↑	F1 ↑	CER ↓	Acc ↑	PESQ ↑	STOI ↑	SI-SDRi ↑	
wav2vec 2.0 BASE [6]	95.04M	5.74	6.43	96.23	0.0233	92.35	88.30	24.77	63.43	2.55	93.9	9.77	64.7
HuBERT BASE [3]	94.68M	5.41	6.42	96.30	0.0736	98.34	88.53	25.20	64.92	2.58	93.9	9.36	65.8
WavLM BASE [5]	94.70M	4.84	6.21	96.79	0.0870	98.63	89.38	22.86	65.94	2.58	94.0	10.37	66.6
C-HuBERT BASE	96.00M	6.14	7.38	96.92	0.0520	97.63	88.95	24.96	65.51	2.63	94.0	11.08	65.7
wav2vec 2.0 LARGE [6]	317.38M	4.75	3.75	96.66	0.0489	95.28	87.11	27.31	65.64	2.52	94.0	10.02	65.5
HuBERT LARGE [3]	316.61M	3.53	3.62	95.29	0.0353	98.76	89.81	21.76	67.62	2.64	94.2	10.45	66.7
WavLM LARGE† [5]	316.62M	3.06	3.44	97.86	0.0886	99.31	92.21	18.36	70.62	2.70	94.5	11.19	68.4
C-HuBERT LARGE	318.95M	3.78	4.02	96.79	0.0329	95.97	89.67	23.16	67.79	2.65	94.3	11.24	66.4

Table 4: Diarization error rate (DER %) results on LibriMix. DER reported for HuBERT and WavLM Base and Large models for 3 speakers and mixture of 2 and 3 speakers were obtained by us. C-HuBERT models are reported with $p_{mix} = 1.0$.

Method	2Mix	3Mix	2Mix + 3Mix
HuBERT BASE [3]	5.88	8.88	9.04
WavLM BASE [5]	4.55	7.13	7.54
C-HuBERT BASE	2.77	4.42	3.95
HuBERT LARGE [3]	5.75	7.84	8.62
WavLM LARGE [5]	3.24	5.77	5.62
C-HuBERT LARGE	2.65	4.08	3.86

clean and test-other when fine-tuned on Libri2Mix instead of LS-10h, lower than the 5.1%/9.1% WER increases for HuBERT.

6.2. Speech Diarization

Table 4 shows results on speech diarization for two-, three-, and a mix of two- and three-speaker datasets. When compared against HuBERT and WavLM BASE and LARGE models, the best DERs across the three settings are attained using C-HuBERT models. In fact, the C-HuBERT BASE model outperforms the WavLM LARGE model on Libri2Mix, Libri3Mix, and Libri(2+3)Mix by 14%, 23%, and 30% relative, respectively. These are impressive gains since the BASE model is considerably smaller than the WavLM LARGE model and was pre-trained on fewer hours of data. Performance on all test sets are further improved when scaling to the LARGE model.

6.3. SUPERB

We compare C-HuBERT with several state-of-the-art SSL models on the SUPERB tasks (Table 3). C-HuBERT shows strong performance on speech enhancement and source separation, which are closer to the pre-training task of C-HuBERT. It lags behind other models on single-speaker tasks such as PR and ASR. However, the performance can be improved and the gap can be reduced by simply scaling C-HuBERT (on PR, 12% PER reduction (1 - 5.41 / 6.14) between HuBERT and C-HuBERT for BASE and 7% gap for LARGE).

6.4. Ablation Studies

We study the effect of mixing probability p_{mix} and max number of speakers K on speech diarization (SD), multi-speaker and single-speaker ASR (MS-ASR and ASR) with the C-HuBERT BASE model (Table 5). On speech diarization, more aggressive mixing (larger K and higher p_{mix}) leads to better results. The trend reverses on single-speaker speech recognition in general. Nevertheless, we observe that C-HuBERT outperforms HuBERT on single-speaker ASR for some

Table 5: Ablation study of mixing parameter: max number of utterances in the mixture K and mixing probability p_{mix} .

K	p_{mix}	SD		MS-ASR			ASR	
		2Mix	3Mix	dev-mix	d-c	d-o	d-c	d-o
2	0.2	3.79	6.70	25.4	4.5	13.1	4.0	9.5
	0.6	3.48	5.96	13.3	6.6	14.9	4.1	9.8
	1.0	3.35	5.78	16.1	5.2	14.4	4.3	10.1
3	0.2	3.65	5.56	26.3	5.0	12.4	4.0	9.2
	0.6	3.33	5.14	18.8	7.1	15.7	4.2	9.5
	1.0	3.01	4.88	17.5	5.1	13.3	4.3	9.8
5	0.2	3.31	5.30	23.5	9.5	18.2	3.9	9.3
	0.6	2.97	4.55	12.7	4.9	13.0	4.4	10.3
	1.0	2.77	4.42	15.6	5.5	13.8	4.6	10.8

configurations (e.g., $K = 5$ and $p = 0.2$ yields 3.9%/9.3% compared to 4.1%/9.4% from HuBERT in Table 1).

We report results on both single- and multi-speaker test sets when fine-tuning C-HuBERT on multi-speaker data (columns below MS-ASR). Overall, $p_{mix} = 0.2$ leads to the worst multi-speaker test results (23.5% to 26.3%), which are still better than those from HuBERT (35.8%). The best multi-speaker test result is obtained with $K = 5$ and $p_{mix} = 0.6$. The results on single-speaker test sets (d-c, d-o) are interesting — with $K = 2$, single-speaker and multi-speaker WERs are negatively correlated, while with $K = 5$ they are positively correlated. We believe the observation arises from the interaction of two factors: how mismatched pre-training and fine-tuning are, and how mismatched pre-training and testing are.

7. CONCLUSION

This paper presents Cocktail HuBERT, a large-scale pre-trained model with the objective of masked pseudo source separation. C-HuBERT extends the HuBERT framework to multiple speakers, enabling the models to outperform the state-of-the-art models on MS-ASR and SD tasks, while achieving competitive performance on other SUPERB single-speaker and multi-speaker tasks.

8. REFERENCES

- [1] Yu-An Chung, Wei-Ning Hsu, Hao Tang, and James Glass, “An unsupervised autoregressive model for speech representation learning,” *arXiv preprint arXiv:1904.03240*, 2019.
- [2] Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli, “wav2vec: Unsupervised pre-training for speech recognition,” *arXiv preprint arXiv:1904.05862*, 2019.
- [3] Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman

- Mohamed, “HuBERT: Self-supervised speech representation learning by masked prediction of hidden units,” *TASLP*, 2021.
- [4] Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatong Gu, and Michael Auli, “Data2vec: A general framework for self-supervised learning in speech, vision and language,” *ICML*, 2022.
- [5] Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioaka, Xiong Xiao, et al., “Wavlm: Large-scale self-supervised pre-training for full stack speech processing,” *STSP*, 2022.
- [6] A. Baevski, Y. Zhou, A., Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *NeurIPS*, 2020.
- [7] Yi-Chen Chen, Po-Han Chi, Shu-wen Yang, Kai-Wei Chang, Jheng-hao Lin, Sung-Feng Huang, Da-Rong Liu, Chi-Liang Liu, Cheng-Kuang Lee, and Hung-yi Lee, “Speechnet: A universal modularized model for speech processing tasks,” *arXiv preprint arXiv:2105.03070*, 2021.
- [8] Wei-Ning Hsu, David Harwath, and James Glass, “Transfer learning from audio-visual grounding to speech recognition,” in *Interspeech*, 2019.
- [9] Ye Jia, Yu Zhang, Ron Weiss, Quan Wang, Jonathan Shen, Fei Ren, Patrick Nguyen, Ruoming Pang, Ignacio Lopez Moreno, Yonghui Wu, et al., “Transfer learning from speaker verification to multispeaker text-to-speech synthesis,” *NeurIPS*, 2018.
- [10] Shu-wen Yang, Po-Han Chi, Yung-Sung Chuang, Cheng-I Jeff Lai, Kushal Lakhota, Yist Y Lin, Andy T Liu, Jiatong Shi, Xuankai Chang, Guan-Ting Lin, et al., “Superb: Speech processing universal performance benchmark,” in *Interspeech*, 2021.
- [11] Hsiang-Sheng Tsai, Heng-Jui Chang, Wen-Chin Huang, Zili Huang, Kushal Lakhota, Shu-wen Yang, Shuyan Dong, Andy T Liu, Cheng-I Jeff Lai, Jiatong Shi, et al., “Superb-sg: Enhanced speech processing universal performance benchmark for semantic and generative capabilities,” in *ACL*, 2022.
- [12] Alexei Baevski, Wei-Ning Hsu, Alexis Conneau, and Michael Auli, “Unsupervised speech recognition,” *NeurIPS*, 2021.
- [13] Alexander H Liu, Cheng-I Jeff Lai, Wei-Ning Hsu, Michael Auli, Alexei Baevski, and James Glass, “Simple and effective unsupervised speech synthesis,” in *Interspeech*, 2022.
- [14] Junrui Ni, Liming Wang, Heting Gao, Kaizhi Qian, Yang Zhang, Shiyu Chang, and Mark Hasegawa-Johnson, “Unsupervised text-to-speech synthesis by unsupervised automatic speech recognition,” in *Interspeech*, 2022.
- [15] Adam Polyak, Yossi Adi, Jade Copet, Eugene Kharonov, Kushal Lakhota, Wei-Ning Hsu, Abdelrahman Mohamed, and Emmanuel Dupoux, “Speech resynthesis from discrete disentangled self-supervised representations,” in *Interspeech*, 2021.
- [16] Kushal Lakhota, Eugene Kharonov, Wei-Ning Hsu, Yossi Adi, Adam Polyak, Benjamin Bolte, Tu-Anh Nguyen, Jade Copet, Alexei Baevski, Abdelrahman Mohamed, et al., “On generative spoken language modeling from raw audio,” *TACL*, 2021.
- [17] Kai-Wei Chang, Wei-Cheng Tseng, Shang-Wen Li, and Hung-yi Lee, “An exploration of prompt tuning on generative spoken language model for speech processing tasks,” in *Interspeech*, 2022.
- [18] Wei-Ning Hsu, Anuroop Sriram, Alexei Baevski, Tatiana Likhomanenko, Qiantong Xu, Vineel Pratap, Jacob Kahn, Ann Lee, Ronan Collobert, Gabriel Synnaeve, et al., “Robust wav2vec 2.0: Analyzing domain shift in self-supervised pre-training,” in *Interspeech*, 2021.
- [19] Kazuya Kawakami, Luyu Wang, Chris Dyer, Phil Blunsom, and Aaron van den Oord, “Learning robust and multilingual speech representations,” *arXiv preprint arXiv:2001.11128*, 2020.
- [20] Alexis Conneau, Alexei Baevski, Ronan Collobert, Abdelrahman Mohamed, and Michael Auli, “Unsupervised cross-lingual representation learning for speech recognition,” in *Interspeech*, 2021.
- [21] Arun Babu, Changan Wang, Andros Tjandra, Kushal Lakhota, Qiantong Xu, Naman Goyal, Kritika Singh, Patrick von Platen, Yatharth Saraf, Juan Pino, et al., “Xls-r: Self-supervised cross-lingual speech representation learning at scale,” in *Interspeech*, 2022.
- [22] Bowen Shi, Wei-Ning Hsu, Kushal Lakhota, and Abdelrahman Mohamed, “Learning audio-visual speech representation by masked multimodal cluster prediction,” in *ICLR*, 2022.
- [23] Wei-Ning Hsu and Bowen Shi, “A single self-supervised model for many speech modalities enables zero-shot modality transfer,” in *NeurIPS*, 2022.
- [24] Jing Shi, Xuankai Chang, Tomoki Hayashi, Yen-Ju Lu, Shinji Watanabe, and Bo Xu, “Discretization and re-synthesis: an alternative method to solve the cocktail party problem,” *arXiv preprint arXiv:2112.09382*, 2021.
- [25] X. Chang, N. Moritz, T. Hori, S. Watanabe, and J. Le Roux, “Extended graph temporal classification for multi-speaker end-to-end asr,” in *ICASSP*, 2022.
- [26] Dong Yu, Morten Kolbæk, Zheng-Hua Tan, and Jesper Jensen, “Permutation invariant training of deep models for speaker-independent multi-talker speech separation,” in *ICASSP*, 2017.
- [27] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *ICASSP*, 2015.
- [28] J. Kahn et al., “Libri-light: A benchmark for asr with limited or no supervision,” in *ICASSP*, 2020.
- [29] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” 2014.
- [30] J. Cosentino, M. Pariente, S. Cornell, A. Deleforge, and E. Vincent, “Librimix: An open-source dataset for generalizable speech separation,” 2020.
- [31] A. Graves, S. Fernandez, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *ICML*, 2006.
- [32] J. Fiscus, J. Ajot, M. Michel, , and J. Garofolo, “The rich transcription 2006 spring meeting recognition evaluation,” in *MLMI*, 2006.
- [33] P. Guo, X. Chang, S. Watanabe, and L. Xie, “Multi-speaker asr combining non-autoregressive conformer ctc and conditional speaker chain,” in *Interspeech*, 2021.
- [34] X. Chang, Y. Qian, K. Yu, and S. Watanabe, “End-to-end monaural multi-speaker asr system without pretraining,” in *ICASSP*, 2019.