

Pac-HuBERT: Self-Supervised Music Source Separation via Primitive Auditory Clustering and Hidden-Unit BERT

Chen, Ke; Wichern, Gordon; Germain, Francois; Le Roux, Jonathan

TR2023-030 May 06, 2023

Abstract

In spite of the progress in music source separation research, the small amount of publicly-available clean source data remains a constant limiting factor for performance. Thus, recent advances in self-supervised learning present a largely-unexplored opportunity for improving separation models by leveraging unlabelled music data. In this paper, we propose a self-supervised learning framework for music source separation inspired by the HuBERT speech representation model. We first investigate the potential impact of the original HuBERT model by inserting an adapted version of it into the well-known Demucs V2 time-domain separation architecture. We then propose Pa-HuBERT, a time-frequency-domain self-supervised model, that we later use in combination with a Res-U-Net decoder for source separation. Pa-HuBERT uses primitive auditory features of music as unsupervised clustering labels to initialize the self-supervised pretraining process using the Free Music Archive (FMA) dataset. The resulting framework achieves better source-to-distortion ratio (SDR) performance on the MusDB18 test set than the original Demucs V2 and Res-U-Net models. We further demonstrate that it can boost performance with small amounts of supervised data. Ultimately, our proposed framework is an effective solution to the challenge of limited clean source data for music source separation.

IEEE ICASSP Satellite Workshop on Self-supervision in Audio, Speech and Beyond (SASB) 2023

Pa \mathcal{G} -HuBERT: SELF-SUPERVISED MUSIC SOURCE SEPARATION VIA PRIMITIVE AUDITORY CLUSTERING AND HIDDEN-UNIT BERT

Ke Chen[✉], Gordon Wichern[✉], François G. Germain[✉], Jonathan Le Roux[✉]

[✉]Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA

[✉]University of California San Diego (UCSD), La Jolla, CA, USA

ABSTRACT

In spite of the progress in music source separation research, the small amount of publicly-available clean source data remains a constant limiting factor for performance. Thus, recent advances in self-supervised learning present a largely-unexplored opportunity for improving separation models by leveraging unlabelled music data. In this paper, we propose a self-supervised learning framework for music source separation inspired by the HuBERT speech representation model. We first investigate the potential impact of the original HuBERT model by inserting an adapted version of it into the well-known Demucs V2 time-domain separation architecture. We then propose Pa \mathcal{G} -HuBERT, a time-frequency-domain self-supervised model, that we later use in combination with a Res-U-Net decoder for source separation. Pa \mathcal{G} -HuBERT uses primitive auditory features of music as unsupervised clustering labels to initialize the self-supervised pretraining process using the Free Music Archive (FMA) dataset. The resulting framework achieves better source-to-distortion ratio (SDR) performance on the MusDB18 test set than the original Demucs V2 and Res-U-Net models. We further demonstrate that it can boost performance with small amounts of supervised data. Ultimately, our proposed framework is an effective solution to the challenge of limited clean source data for music source separation.

Index Terms— Music source separation, primitive auditory principles, self-supervised Learning, BERT

1. INTRODUCTION

Music source separation aims to separate one or more sound sources (e.g., vocals, drums, bass, and other instruments) from music tracks. This task has broad applications in various domains, including vocal-accompaniment separation and music remixing. Current state-of-the-art methods for the task rely on deep learning techniques [1]. However, unlike in speech research [2], data for music source separation remains limited. The benchmark dataset for the task, MusDB18 [3], consists of only 100 songs for training and 50 songs for testing. The scarcity of data is not only due to the challenge of collecting source data, but also to copyright issues of music assets.

Self-supervised learning (SSL) is a rapidly growing deep learning approach that leverages unlabeled data to improve model training, resulting in better generalization capabilities and task performance [4]. SSL has already achieved promising results in a multitude of audio processing tasks such as speech recognition [5], speech quality prediction [6, 7], or audio classification [8–10], and has even demonstrated the ability to generalize across tasks [11, 12]. Recently, several SSL approaches for music representation learning have also

been proposed [13–15]. However, all of these approaches focus on the learning of representations for content labeling tasks, while source separation requires a representation capable of isolating individual components present in an audio mixture.

One area where SSL has been used for source separation is using the pretrained WavLM [16] model for speech enhancement [17, 18]. WavLM is based on HuBERT [5], one of the most effective speech SSL models. HuBERT leverages a transformer-based BERT architecture [19] and is trained to generate a discrete token at each time step, using a masking mechanism such that HuBERT learns to predict the masked tokens from the input data and the unmasked tokens. The tokens are cluster indices, initially determined by running k-means clustering on MFCC features of unlabeled speech data, and subsequent HuBERT iterations use tokens obtained by clustering intermediate network layer outputs from earlier iterations.

In this paper, we explore the use of HuBERT-style methods for enhancing the quality of music source separation models by using readily available unlabeled (i.e., containing no isolated sources or stems) music mixtures. However, the original HuBERT may not be directly applicable to music separation due to its lower 16 kHz sampling rate, and the MFCC features used for creating HuBERT training targets may not provide enough time-frequency (TF) resolution to provide a useful representation for separating music sources. Therefore, we explore replacing MFCC features in the initial step of HuBERT with TF features specifically designed for music separation such as common-fate (2DFT) [20], repetition structure (REPET) [21], harmonic-percussive source separation (HPSS) [22], and melody contour (Melodia) [23]. These primitive auditory features, which in some sense mimic how the human brain performs an initial segmentation of an auditory scene [24], have previously been used to create pseudo-labels for unsupervised music source separation [25], but never in a HuBERT-style training paradigm.

One fundamental design decision for both SSL and source separation models is whether they operate on time-domain waveforms or TF representations such as spectrograms. Time-domain audio representation models such as HuBERT and wav2vec [26] encode raw audio signals into a latent representation, while time-domain separation models contain similar encoders, but also include decoders to obtain separated signals, either through direct synthesis (e.g., Wave-U-Net [27]) or through masking of features (e.g., TasNet [28]). These models are widely used in speech separation, where they offer high separation quality and low latency in real-time scenarios. However, non-speech sources such as music and general sounds are typically recorded at a higher sampling rate and contain a wide variety of complex harmonic patterns, which is why TF-domain models (or hybrid models) have been found historically to yield better performance in audio classification [8, 29] and music source separation [30–34].

We introduce a novel framework that leverages self-supervised learning in either the time domain or TF domain. Our main con-

This work was performed while Ke Chen was an intern at MERL.

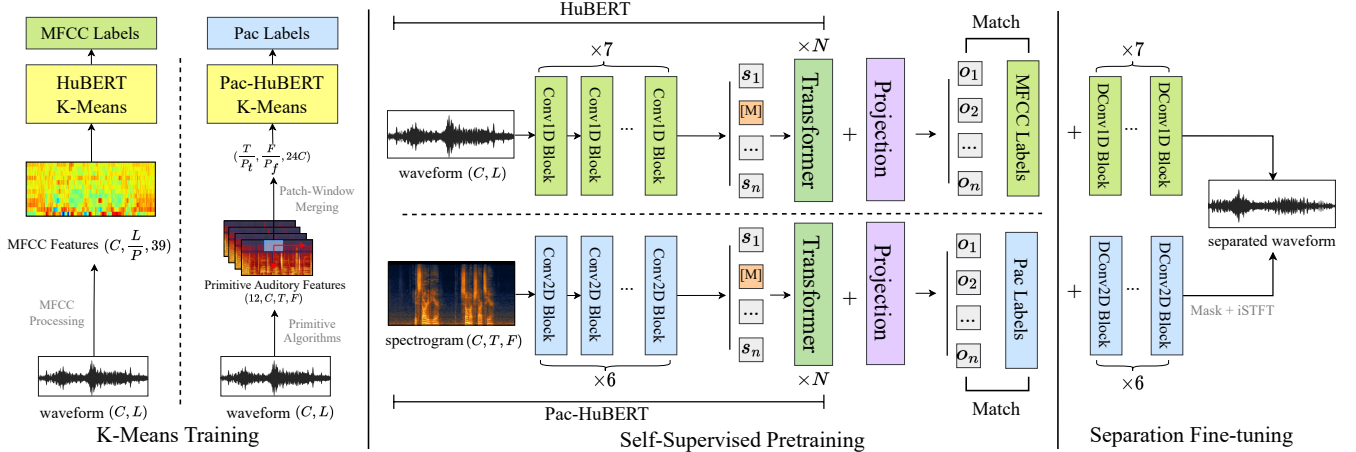


Fig. 1: The whole pipeline of the proposed self-supervised music source separation framework. Left: k-means models for label creation; Middle: self-supervised pretraining (HuBERT and our proposed PaG-HuBERT); Right: fine-tuning on separation tasks.

tributions are as follows. First, we adapt HuBERT into a time-domain music source separation model, which provides support for subsequent model design in the TF domain. Next, we introduce the first self-supervised framework for TF-domain music source separation, PaG-HuBERT (pronounced Pac-HuBERT, for primitive auditory clustering HuBERT). Our framework leverages primitive auditory features of music data to initiate the self-supervised pretraining process and embeds the resulting pretrained layers into a Res-U-Net [31] separation model, an already very effective architecture for the task. Ultimately, we find the separation performance of models including either our time-domain or TF-domain pretrained layers to outperform their original counterparts. Additionally, we evaluate performance in scenarios with even less available training data, demonstrating the robustness of the pretrained model representation, with reliable performance with only a small amount of data available for supervised fine-tuning.

2. MODEL ARCHITECTURE

Our proposed self-supervised music source separation framework (see Fig. 1) is composed of an encoder (convolution blocks), a bottleneck model (transformer), and a decoder (deconvolution blocks). The learning pipeline consists of k-means training, self-supervised pretraining, and separation fine-tuning.

2.1. Encoder Model

The encoder acts as a feature extractor for audio waveforms or spectrograms, as depicted in the middle of Fig. 1. For the time-domain model at the top, the encoder consists of 7 downsampling 1D convolutional layer blocks (CNN). The input audio waveform, denoted by $x \in \mathbb{R}^{C \times L}$, where C denotes the number of channels and L the sample length, is fed into the encoder and transformed into the bottleneck feature $s \in \mathbb{R}^{C_b \times \frac{L}{P}}$, with C_b the channel size and P the downsampling rate. We follow the design of HuBERT to construct the 1D CNN blocks, each containing one CNN layer and one GeLU function and also similar to the encoder block in Demucs V2 [35] (see Fig. 2). In our experiments, in order to use the pretrained model of HuBERT while still following the training scheme of Demucs V2, each input is a 3-second mono audio signal sampled at 16 kHz ($C = 1$, $L = 48000$), and we set $C_b = 1024$ and $P = 320$.

For the TF-domain model at the bottom, the encoder consists of

6 downsampling 2D CNN blocks. The input audio signals are first transformed into STFT spectrograms, denoted as $x \in \mathbb{R}^{C \times T \times F}$, where C denotes the number of channels, T the number of frames, and F the number of frequencies. The encoder then converts x into the bottleneck feature $s \in \mathbb{R}^{C_b \times \frac{T}{P_t} \times \frac{F}{P_f}}$, where P_t and P_f denote the downsampling rates on the time axis and the frequency axis, respectively. We follow the Res-U-Net [31] separation model for the encoder block design (see Fig. 2 for details). Since we build our own TF-domain SSL model, we can align it with the traditional music source separation pipeline. Hence, each input is a 3-second stereo audio sampled at 44.1 kHz ($C = 2$, $L = 132300$). The STFT window size is 2048 and the hop size is 441, resulting in $T = 320$ frames after zero-padding with 20 frames to the right, and $F = 1024$ frequency bins after removing the Nyquist frequency. The model parameters are set to $C_b = 384$, $P_t = 32$, and $P_f = 64$.

2.2. Bottleneck Model

The bottleneck model consists of N transformer encoder blocks, each of which contains a multi-head self-attention layer [19] with d heads, hidden size h , and a feedforward layer of inner hidden size $4h$. For time-domain models, we use the HuBERT-LARGE model ($N = 24$, $d = 16$, $h = 1024$). For TF-domain models, to limit GPU memory consumption and without existing prior pretrained models, we set $N = 12$, $d = 8$, and $h = 384$. The output of the encoder s is reshaped to $(\frac{L}{P}, C_b)$ in the time-domain case and $(\frac{T}{P_t} \times \frac{F}{P_f}, C_b)$ in the TF-domain case, and fed into the bottleneck model. The output is a latent feature s' with the same shape as s , which is sent to the decoder for the separation task (Section 2.3) or the projection module for self-supervised pretraining (Section 2.4). We denote the encoder/bottleneck combinations as HuBERT for the time-domain model and as PaG-HuBERT for the TF-domain model.

2.3. Decoder Model

The decoder is designed by replacing all CNN layers in the encoder with deconvolutional layers (DCNN) for both time-domain and TF-domain models (see right of Fig. 1). The upsampling rate is equal to the downsampling rate in order for the audio input and the separation output (or mask) to have the same shape. Additionally, we have skip-connections between each encoder block and the corresponding decoder block (not shown in Fig. 1).

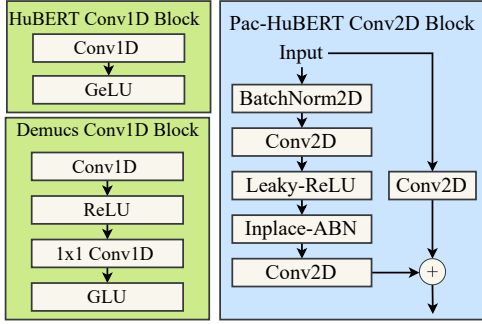


Fig. 2: Comparison of CNN block designs between Demucs V2 [35], HuBERT [5], and PaG-HuBERT (the same as Res-U-Net [31]).

2.4. Training and Inference Scheme

K-means Training: HuBERT creates initial labels for SSL by performing k-means clustering on MFCC features (see left of Fig. 1) as follows: 1) extract 39-dimensional MFCC features with 20 ms frame size and 10 ms frame shift, from the LibriSpeech and Libri-Light audio samples [36, 37], and 2) fit a 500-cluster k-means model to the MFCC features, and use the cluster index assigned to each time frame as the label.

For PaG-HuBERT, we replace MFCC features for each time frame with primitive auditory features computed over spectrogram patches (see left of Fig. 1). Formally, given a mixture spectrogram $x \in \mathbb{R}^{C \times T \times F}$, each of the primitive algorithms we use here generates foreground and background estimates $e_f, e_b \in \mathbb{R}^{C \times T \times F}$. The six primitive algorithms we employ¹ are HPSS [22, 39], REPET [21], REPET-SIM [40], FT2D-M, FT2D-R [20], and Melodia [23]. Combining foreground and background cues for each primitive algorithms, we obtain a 12-dimensional feature for each TF bin. Due to the memory requirements of transformers, we follow the audio spectrogram transformer [8], combining features in non-overlapping spectrogram patches of shape (P_t, P_f) bins (matching the downsampling ratio in the TF-domain encoder). We obtain a 24-dimensional feature for each patch by first dividing each patch into a low-frequency sub-patch containing the lower $P_f/2$ TF bins, and a high-frequency sub-patch containing the upper $P_f/2$ TF bins. We then average over TF bins in each sub-patch, obtaining two 12-dimensional features, which we concatenate. For stereo ($C = 2$) music data, we obtain a 48-dimensional feature vector for each TF patch, which is of comparable dimension to the 39-dimensional MFCCs used in the original HuBERT model. We fit a 960-cluster k-means model to all of the 48-dimensional features extracted from the FMA-Large dataset [41] (106 574 music tracks or about 890 h).

SSL Pretraining via Masked Unit Prediction: We train HuBERT and PaG-HuBERT with the k-means labels using masking prediction [42] for the self-supervised learning process. The bottleneck model of HuBERT (or PaG-HuBERT) produces a sequence of n labels with D classes. For HuBERT, $n = \frac{T}{P_t}$ and $D = 500$; for PaG-HuBERT, $n = (\frac{T}{P_t} \times \frac{F}{P_f})$ and $D = 960$. We connect the bottleneck model to a projection module (see middle of Fig. 1) to map the output feature into the class feature $o \in \mathbb{R}^{n \times E}$, where E is the projection dimension. The loss function is defined as:

$$L = \frac{1}{|M|} \sum_{t \in M} \frac{\exp(\tau(o_t \cdot e_c) / (\|o_t\| \|e_c\|))}{\sum_{c'=1}^D \exp(\tau(o_t \cdot e_{c'}) / (\|o_t\| \|e_{c'}\|))}, \quad (1)$$

¹Implementations from <https://github.com/nussl/nussl> [38]

where M is the set containing all indices of masked tokens (as shown in Fig. 1), o_t is the class feature token at frame t in HuBERT (or TF patch t in PaG-HuBERT), e_c is a learned embedding for k-means cluster c to which the frame t was assigned, and $\tau = 10$ scales the logit. We apply the span masking mechanism by randomly selecting $p\%$ of the frames as starting indices and masking spans of l frames. For PaG-HuBERT, $p = 40\%$ and $l = 5$. The loss is only computed over the masked frames.

Source Separation Fine-tuning: After pretraining the encoder and the bottleneck models with LibriSpeech+Libri-Light or FMA-Large, we connect the bottleneck model to the decoder and fine-tune the whole model on the separation task with separation datasets. The loss function is L1 Loss between the final separation signals and the source signals, the same as both Demucs V2 and Res-U-Net.

3. EXPERIMENTS

3.1. Hyperparameters and Training Details

During the pretraining stage, for HuBERT, we directly used the pretrained model HuBERT-LARGE from the torchaudio² library. HuBERT-LARGE is pretrained on the Libri-Light speech dataset. The details can be found in [5]. For PaG-HuBERT, we pretrained it on the ‘‘FMA-Large’’ set of the Free Music Archive dataset, containing 890 hours of music tracks. We applied the AdamW optimizer with a 32 000-step warm-up [43] and a reduced LR scheduler after 150 000 steps. The basic learning rate was 5×10^{-4} . We trained the model using a batch size of 96 on 8 NVIDIA A40 GPUs for 180 000 steps, by which time the masked accuracy on the validation set achieves 81%. Since higher masked accuracy does not guarantee better separation performance, we used pretrained checkpoints at various steps and fine-tuned them on a subset of the separation training set to determine the best model (details in Table 3). The checkpoint after 60 000 steps yielded the best separation performance and we used it for all subsequent experiments.

During the fine-tuning stage, we used the MusDB18 [3] dataset for training separation models. We followed [35] to divide the 100-song training set into 84 songs for training and 16 songs for validation. After determining the best model on the validation set, we evaluate it on the test set of 50 songs. For the time-domain framework, we trained the decoder while fine-tuning the HuBERT-LARGE model. For the TF-domain framework, we trained the decoder while fine-tuning the PaG-HuBERT model. The data formats and audio processing settings are mentioned in Section 2.1. We applied the Adam optimizer with a learning rate of 3×10^{-4} (time domain) or 10^{-3} (TF domain). We adopted a 3000-step warm up and a decay scheduler, where the learning rate is scaled by $\alpha = 0.9$ every 15 000 steps. We trained the models using a batch size of 128 (time domain) or 96 (TF domain) on 8 NVIDIA A40 GPUs, for 200 000 steps, by which time all models had converged in terms of the separation performance on the validation set.

3.2. Separation Results

Effectiveness of Self-Supervised Learning: Table 1 presents the source-to-distortion ratio (SDR) performance on MusDB18. We follow the MusDB18 benchmark using the SiSEC2018 [46] version of the SDR metric (BSS Eval v4 framewise SDR) implemented by `mus_eval`³. We report the median SDR over all 50 songs in the MusDB18 test set.

²<https://github.com/pytorch/audio>

³<https://github.com/sigsep/sigsep-mus-eval>

Table 1: Source-to-distortion ratio (SDR) performance on the test set of MusDB18. Results marked in grey are not directly comparable due to the use of additional training data or the absence of a validation set. All results are at 44.1 kHz unless otherwise noted.

Model	Pretrain	SDR on MusDB18 test set (dB)			
		Vocals	Drums	Bass	Other
Demucs V2 [35] - 16 kHz	✗	5.02	6.02	5.40	3.41
HuBERT-SEP ^N - 16 kHz	✗	5.14	5.59	4.96	2.82
HuBERT-SEP ^{SSL} - 16 kHz	✓	5.58	6.63	6.03	3.65
Bytedance Res-U-Net [31]	✗	7.83	5.47	5.21	4.90
PaG-HuBERT-SEP ^N	✗	8.07	5.78	5.21	5.29
PaG-HuBERT-SEP ^{SSL}	✓	8.32	5.86	6.01	5.38
PaG-HuBERT-SEP ^{2SSL}	✓	8.52	6.20	5.76	5.18
Open Unmix [44]	✗	6.32	5.73	5.23	4.02
Spleeter [34]	✗	6.86	6.71	5.51	4.55
D3Net [45]	✗	7.24	6.68	5.25	4.53
MDX-Net [30]	✗	9.00	7.33	7.86	5.95
Band-Split RNN [32]	✗	10.01	9.01	7.22	6.70
HT Demucs [33]	✗	7.93	7.94	8.48	5.72

Table 2: Separation performance with different amounts of training data. Subscripts denote improvement from SSL pretraining.

Model	Data Ratio	SDR on MusDB18 test set (dB)			
		Vocals	Drums	Bass	Other
PaG-HuBERT-SEP ^N	25%	5.72	4.42	4.00	3.38
	50%	6.76	4.50	4.14	4.13
	100%	8.07	5.78	5.21	5.29
PaG-HuBERT-SEP ^{SSL}	25%	6.75 _{+1.03}	4.57 _{+0.15}	4.11 _{+0.11}	3.90 _{+0.52}
	50%	7.32 _{+0.56}	5.16 _{+0.66}	5.07 _{+0.93}	4.50 _{+0.37}
	100%	8.32 _{+0.25}	5.86 _{+0.08}	6.01 _{+0.80}	5.38 _{+0.09}

We denote our models as *HuBERT-SEP* and *PaG-HuBERT-SEP*, with a superscript indicating no pretraining (*N*) or SSL pretraining (*SSL*). Similar to HuBERT, we clustered the latent features of the 6th transformer block from the first pretrained model, and pretrained a new PaG-HuBERT as the second pretraining iteration, indicated as *2SSL*. All experiments on the original Res-U-Net and Demucs V2 were reproduced by us to ensure a fair comparison. Time-domain models were trained at 16 kHz to utilize the pretrained HuBERT model, so the baseline Demucs V2 results differ from [35]. Additionally, we provide the reported performance of several state-of-the-art models, although some of them (denoted in grey) used additional training data or were developed for the MDX challenge [1], which used the MusDB18 test set as a validation set for model selection.

From the comparison between Demucs V2, Res-U-Net, HuBERT-SEP^N, and PaG-HuBERT^N models (i.e., without pretraining), we observe that PaG-HuBERT^N generally outperforms Res-U-Net. The potential reason is that the transformer bottleneck processes the encoder features more efficiently than the CNN bottleneck model in Res-U-Net. However, we found that Demucs V2 performs better than HuBERT-SEP^N, possibly due to the more complex encoder-decoder design of Demucs V2, which was also reported in [35].

When fine-tuning the pretrained HuBERT and PaG-HuBERT models on the separation task, we observe an improvement in SDR for all four sources compared to models without pretraining. This demonstrates the effectiveness of HuBERT in contributing to audio separation tasks beyond speech, and also shows that our proposed PaG-HuBERT model, pretrained with primitive auditory features from music data, can improve the performance of TF-domain

Table 3: Performance of pretrained models at different training steps on the pretraining and separation tasks.

Metrics	Training Steps			
	30 000	60 000	120 000	180 000
Masked Accuracy (Valid.)	0.71	0.78	0.79	0.81
Vocal SDR (25% Data)	6.02	6.75	6.50	6.04

separation models. When comparing PaG-HuBERT-SEP^{SSL} and PaG-HuBERT-SEP^{2SSL}, we observe that *2SSL* yields better performance on vocals and drums separation, but is degraded on bass and other separation. We believe that the design of iterative pretraining can be further improved, including exploring different layers of latent features for pretraining and training varying numbers of k-means models for different separation sources. This will be a focus of our future work. Specifically, HuBERT shows a larger improvement than PaG-HuBERT, which can potentially be attributed to the size of the pretraining dataset (60 000 hours for HuBERT vs. 890 hours for PaG-HuBERT). The ability of speech-only pretraining to improve music separation performance presents another opportunity for future work in finding ways to harness this data for full-bandwidth TF-domain music separation models.

Effectiveness with Limited Training Data: We further extracted from the 84 MusDB18 training songs a 25% data subset with 21 songs, and a 50% data subset with 42 songs. We then fine-tuned our PaG-HuBERT-SEP model on these subsets and report the separation performance of the test set in Table 2. Results show that the pretrained model achieves significantly better performance than the models without pretraining on both subsets. Notably, with only 25% of the supervised data, the pretrained model improves separation performance for vocals and other source classes by 0.5 to 1.0 dB. This demonstrates the effectiveness of pretraining using primitive auditory features, as the foreground and background features are highly correlated with the vocals and other source classes. However, for bass and drums, pretraining provides little benefit with 25% of the supervised data, indicating that pretraining is less useful for these classes, and different auditory primitives should be explored.

Pretrained Model Selection: We also investigated varying the number of PaG-HuBERT pretraining steps and evaluating vocal separation performance when fine-tuning on 25% of the MusDB18 training data in Table 3. We observe that the separation performance of the model is not consistent with its pretraining performance (i.e., masked token prediction accuracy). A similar trend was reported in [9] for audio classification. High accuracy during pretraining may indicate over-fitting, which could weaken the quality of the learned representation and negatively impact performance in separation tasks.

4. CONCLUSION

In this paper, we investigated the efficacy of self-supervised learning methods in music source separation. First, we adapted HuBERT into a time-domain source separation model and achieved notable improvements. Next, we proposed PaG-HuBERT, a TF-domain self-supervised source separation model that leverages the primitive auditory features of unlabelled music data during pretraining. PaG-HuBERT demonstrates significant improvements in TF-domain music source separation models, with consistent performance increases across different proportions of supervised data. We consider PaG-HuBERT to be an effective solution for utilizing unlabeled music data in music source separation. Moving forward, we plan to adapt PaG-HuBERT to other separation models besides Res-U-Net, and explore its generalization on the separation of more instruments.

5. REFERENCES

- [1] Y. Mitsufuji, G. Fabbro, S. Uhlich, F.-R. Stöter, A. Défossez, M. Kim, W. Choi, C.-Y. Yu, and K.-W. Cheuk, “Music demixing challenge 2021,” *Front. Signal Process.*, vol. 1, 2022.
- [2] R. Haeb-Umbach, S. Watanabe, T. Nakatani, M. Bacchiani, B. Hoffmeister, M. L. Seltzer, H. Zen, and M. Souden, “Speech processing for digital home assistants: Combining signal processing with deep-learning techniques,” *IEEE Signal Process. Mag.*, vol. 36, no. 6, pp. 111–124, 2019.
- [3] Z. Rafii, A. Liutkus, F.-R. Stöter, S. I. Mimilakis, and R. Bittner, “The MUSDB18 corpus for music separation,” 2017.
- [4] L. Ericsson, H. Gouk, C. C. Loy, and T. M. Hospedales, “Self-supervised representation learning: Introduction, advances, and challenges,” *IEEE Signal Process. Mag.*, vol. 39, no. 3, pp. 42–62, 2022.
- [5] 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 Trans. Audio, Speech, Lang. Process.*, vol. 29, pp. 3451–3460, 2021.
- [6] W.-C. Tseng, C.-Y. Huang, W.-T. Kao, Y. Y. Lin, and H.-Y. Lee, “Utilizing self-supervised representations for MOS prediction,” in *Proc. Interspeech*, 2021, pp. 2781–2785.
- [7] H. Becerra, A. Ragano, and A. Hines, “Exploring the influence of fine-tuning data on wav2vec 2.0 model for blind speech quality prediction,” in *Proc. Interspeech*, 2022, pp. 4088–4092.
- [8] Y. Gong, C.-I. Lai, Y.-A. Chung, and J. Glass, “SSAST: Self-supervised audio spectrogram transformer,” in *Proc. AAAI*, vol. 36, no. 10, 2022, pp. 10 699–10 709.
- [9] P.-Y. Huang, H. Xu, J. B. Li, A. Baevski, M. Auli, W. Galuba, F. Metze, and C. Feichtenhofer, “Masked autoencoders that listen,” in *Proc. NeurIPS*, 2022.
- [10] S. Chen, Y. Wu, C. Wang, S. Liu, D. Tompkins, Z. Chen, and F. Wei, “BEATs: Audio pre-training with acoustic tokenizers,” *arXiv preprint arXiv:2212.09058*, 2022.
- [11] J. Turian, J. Shier, H. R. Khan, B. Raj, B. W. Schuller, C. J. Steinmetz, C. Malloy, G. Tzanetakis, G. Velarde, K. McNally *et al.*, “HEAR: Holistic evaluation of audio representations,” in *Proc. 2021 Compet. Demonstrations Track (NeurIPS)*, 2022, pp. 125–145.
- [12] H.-H. Wu, P. Seetharaman, K. Kumar, and J. P. Bello, “Wav2clip: Learning robust audio representations from clip,” in *Proc. ICASSP*, 2022, pp. 4563–4567.
- [13] H. Zhu, Y. Niu, D. Fu, and H. Wang, “MusicBERT: A self-supervised learning of music representation,” in *Proc. ACM Multimed.*, 2021, pp. 3955–3963.
- [14] H.-H. Wu, C.-C. Kao, Q. Tang, M. Sun, B. McFee, J. P. Bello, and C. Wang, “Multi-task self-supervised pre-training for music classification,” in *Proc. ICASSP*, 2021, pp. 556–560.
- [15] A. N. Carr, Q. Berthet, M. Blondel, O. Teboul, and N. Zeghidour, “Self-supervised learning of audio representations from permutations with differentiable ranking,” *IEEE Signal Process. Lett.*, vol. 28, pp. 708–712, 2021.
- [16] 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 J. Sel. Top. Signal Process.*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [17] H. Song, S. Chen, Z. Chen, Y. Wu, T. Yoshioka, M. Tang, J. W. Shin, and S. Liu, “Exploring WavLM on speech enhancement,” in *Proc. SLT*, 2022, pp. 451–457.
- [18] X.-Y. Zhao, Q.-S. Zhu, and J. Zhang, “Speech enhancement using self-supervised pre-trained model and vector quantization,” in *Proc. AP-SIPA ASC*, 2022, pp. 330–334.
- [19] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Proc. NeurIPS*, 2017.
- [20] P. Seetharaman, F. Pishdadian, and B. Pardo, “Music/voice separation using the 2D fourier transform,” in *Proc. WASPAA*, 2017, pp. 36–40.
- [21] Z. Rafii and B. Pardo, “Repeating pattern extraction technique (REPET): A simple method for music/voice separation,” *IEEE Trans. Speech Audio Process.*, vol. 21, no. 1, pp. 71–82, 2013.
- [22] J. Driedger, M. Müller, and S. Disch, “Extending harmonic-percussive separation of audio signals,” in *Proc. ISMIR*, 2014, pp. 611–616.
- [23] J. Salamon and E. Gómez, “Melody extraction from polyphonic music signals using pitch contour characteristics,” *IEEE Trans. Speech Audio Process.*, vol. 20, no. 6, pp. 1759–1770, 2012.
- [24] J. H. McDermott, D. P. W. Ellis, and E. P. Simoncelli, “Empirical derivation of acoustic grouping cues from natural sound statistics,” in *Proc. Assoc. Res. Otolaryngol., Annu. Meet.*, 2011.
- [25] P. Seetharaman, G. Wichern, J. Le Roux, and B. Pardo, “Bootstrapping unsupervised deep music separation from primitive auditory grouping principles,” in *Proc. Workshop Self-Superv. Audio Speech*, 2020.
- [26] S. Schneider, A. Baevski, R. Collobert, and M. Auli, “wav2vec: Unsupervised pre-training for speech recognition,” in *Proc. Interspeech*, 2019, pp. 3465–3469.
- [27] D. Stoller, S. Ewert, and S. Dixon, “Wave-U-Net: A multi-scale neural network for end-to-end audio source separation,” in *Proc. ISMIR*, 2018, pp. 334–340.
- [28] Y. Luo and N. Mesgarani, “TaSNet: Time-domain audio separation network for real-time, single-channel speech separation,” in *Proc. ICASSP*, 2018, pp. 696–700.
- [29] K. Chen, X. Du, B. Zhu, Z. Ma, T. Berg-Kirkpatrick, and S. Dubnov, “HTS-AT: A hierarchical token-semantic audio transformer for sound classification and detection,” in *Proc. ICASSP*, 2022, pp. 646–650.
- [30] M. Kim, W. Choi, J. Chung, D. Lee, and S. Jung, “KUIELab-MDX-Net: A two-stream neural network for music demixing,” in *Proc. MDX Workshop*, 2021, pp. 1–7.
- [31] Q. Kong, Y. Cao, H. Liu, K. Choi, and Y. Wang, “Decoupling magnitude and phase estimation with deep resunet for music source separation,” in *Proc. ISMIR*, 2021, pp. 342–349.
- [32] Y. Luo and J. Yu, “Music source separation with band-split RNN,” *CoRR*, vol. abs/2209.15174, 2022.
- [33] S. Rouard, F. Massa, and A. Défossez, “Hybrid transformers for music source separation,” *arXiv preprint arXiv:2211.08553*, 2022.
- [34] R. Hennequin, A. Khlif, F. Voituret, and M. Moussallam, “Spleeter: a fast and efficient music source separation tool with pre-trained models,” *J. Open Source Softw.*, vol. 5, no. 50, p. 2154, 2020.
- [35] A. Défossez, N. Usunier, L. Bottou, and F. R. Bach, “Music source separation in the waveform domain,” *CoRR*, vol. abs/1911.13254, 2019.
- [36] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “LibriSpeech: An ASR corpus based on public domain audio books,” in *Proc. ICASSP*, 2015, pp. 5206–5210.
- [37] J. Kahn, M. Rivière, W. Zheng, E. Kharitonov, Q. Xu, P. Mazaré, J. Karadayi, V. Liptchinsky, R. Collobert, C. Fuegen *et al.*, “Libri-Light: A benchmark for ASR with limited or no supervision,” in *Proc. ICASSP*, 2020, pp. 7669–7673.
- [38] E. Manilow, P. Seetharaman, and B. Pardo, “The Northwestern University source separation library,” in *Proc. ISMIR*, 2018.
- [39] N. Ono, K. Miyamoto, J. Le Roux, H. Kameoka, and S. Sagayama, “Separation of a monaural audio signal into harmonic/percussive components by complementary diffusion on spectrogram,” in *Proc. EU-SIPCO*, 2008.
- [40] Z. Rafii and B. Pardo, “Music/voice separation using the similarity matrix,” in *Proc. ISMIR*, 2012, pp. 583–588.
- [41] M. Defferrard, K. Benzi, P. Vandergheynst, and X. Bresson, “FMA: A dataset for music analysis,” in *Proc. ISMIR*, 2017, pp. 316–323.
- [42] J. Devlin, M. Chang, K. Lee, and K. Toutanova, “BERT: pre-training of deep bidirectional transformers for language understanding,” in *Proc. NAACL-HLT 2019*, pp. 4171–4186.
- [43] P. Goyal, P. Dollár, R. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, A. Tulloch, Y. Jia, and K. He, “Accurate, large minibatch SGD: Training ImageNet in 1 hour,” *arXiv preprint arXiv:1706.02677*, 2017.
- [44] F. Stöter, S. Uhlich, A. Liutkus, and Y. Mitsufuji, “Open-Unmix - A reference implementation for music source separation,” *J. Open Source Softw.*, vol. 4, no. 41, p. 1667, 2019.
- [45] N. Takahashi and Y. Mitsufuji, “D3Net: Densely connected multilayered densenet for music source separation,” *arXiv preprint arXiv:2211.08553*, 2021.
- [46] F.-R. Stöter, A. Liutkus, and N. Ito, “The 2018 signal separation evaluation campaign,” in *Proc. LVA/ICA*, 2018, pp. 293–305.