

# UniSep: Universal Target Audio Separation with Language Models at Scale

Yuan Yuan Wang<sup>1</sup>, Hangting Chen<sup>2</sup>, Dongchao Yang<sup>1</sup>, Weiqin Li<sup>3</sup>, Dan Luo<sup>3</sup>,  
Guangzhi Li<sup>2</sup>, Shan Yang<sup>2</sup>, Zhiyong Wu<sup>1,3</sup>, Helen Meng<sup>1</sup>, Xixin Wu<sup>1,†</sup>

<sup>1</sup> The Chinese University of Hong Kong, Hong Kong SAR, China

<sup>2</sup> Tencent AI Lab, Audio and Speech Signal Processing Oteam, China

<sup>3</sup> Shenzhen International Graduate School, Tsinghua University, Shenzhen, China

**Abstract**—We propose Universal target audio Separation (UniSep), addressing the separation task on arbitrary mixtures of different types of audio. Distinguished from previous studies, UniSep is performed on unlimited source domains and unlimited source numbers. We formulate the separation task as a sequence-to-sequence problem, and a large language model (LLM) is used to model the audio sequence in the discrete latent space, leveraging the power of LLM in handling complex mixture audios with large-scale data. Moreover, a novel pre-training strategy is proposed to utilize audio-only data, which reduces the efforts of large-scale data simulation and enhances the ability of LLMs to understand the consistency and correlation of information within audio sequences. We also demonstrate the effectiveness of scaling datasets in an audio separation task: we use large-scale data (36.5k hours), including speech, music, and sound, to train a universal target audio separation model that is not limited to a specific domain. Experiments show that UniSep achieves competitive subjective and objective evaluation results compared with single-task models.

**Index Terms**—universal target audio separation, language models, large-scale data

## I. INTRODUCTION

Audio source separation separates individual sources within an audio mixture [1], [2]. In recent years, audio source separation has shown impressive performance in various separation tasks, such as speech separation [3], [4], sound separation [5]–[8], and music source separation [9], [10]. The majority of previous studies focus on source-specific tasks, such as separating distinct speakers in speech separation and isolating different music sources in music separation. To accomplish these source-specific tasks, models specific to each source must be trained on datasets designed for the respective task [11].

In contrast to source-specific separation tasks, some research has proposed universal source separation [5], [12], [13]. One such approach is target audio source separation, which aims to isolate target audio from mixtures. For instance, chen *et al.* [13] use weakly-labeled data and query-based learning for target audio separation, but their focus remains solely on sound. Most existing other methods for

universal source separation depend on supervised learning techniques and utilize limited training datasets [5], [6], [12]. Considering the extensive variety of sounds in daily life, these datasets are limited to accurately represent the majority of audio sources [11]. Recently, GASS [11] utilizes large-scale supervised datasets to train a model that is capable of separating various types of audio in blind source separation tasks. However, a universal solution for target audio separation has not yet been explored. Most approaches typically focus on separating specific domains and sources. They are unable to simultaneously separate a broader range of types, including speech, music, and sound events.

In this study, we propose to build a universal target audio separation model (UniSep) performed on unlimited source domains and unlimited source numbers, which meets the following research problems: (1) how to model the complex audio signal in one unified model, especially, different types of audio commonly have distinct patterns, such as frequency span. (2) whether a unified target audio separation is promising and beneficial compared to previous specific models. (3) whether large-scale audio data can be used to improve the performance of the target audio separation task. We will answer these questions in this study. The main contributions of this study are summarized as follows:

- Instead of modeling the audio data in spectrogram or waveform, we use a pre-trained audio codec model [14] to convert the audio signal into discrete tokens and model the audio data in a discrete latent space. By using language models with large-scale parameters [15], [16], we formulate the traditional audio separation problem as sequence-to-sequence modeling, improving the simplicity of audio separation.
- We propose a novel pre-training method that utilizes large-scale audio-only data, reducing costly data simulation requirements. This method aims to enhance the ability of the LLMs to understand the consistency and correlation of information within audio sequences, which builds the foundation for the separation tasks.
- In addition, we scale up the training dataset to 36.5k hours. This enables us to harness the language model’s power to handle various source types, ultimately achiev-

Demo can be found in <https://uniseparation.github.io/UniSep/>.

<sup>†</sup>Corresponding author. This work is partially supported by the General Research Fund from the Research Grants Council of Hong Kong SAR Government (Project No. 14202623) and the Center for Perceptual and Interactive Intelligence (CPII) Ltd under the Innovation and Technology Commission’s InnoHK Scheme.

ing universal target audio separation.

- To the best of our knowledge, we offer the first study on LLM-based universal audio separation. Experimental results show that the UniSep model is not only capable of performing the universal separation task effectively but can also be fine-tuned with a minimal amount of data to accomplish other related tasks, *e.g.* language-queried audio source separation task.

## II. METHOD

In this section, we present the details of UniSep. We first introduce how to quantize the audio signal into discrete tokens, then we introduce how to use a language model for target audio separation. Lastly, we introduce the proposed pre-training strategy.

### A. Audio Tokenization

In this work, we try to separate any type of audio (speech, sound, music) in a unified discrete latent space. To obtain such latent space, we opt for neural audio codec models [14]. Specifically, Assuming an audio signal of duration  $d$  with sample rate  $f_s$  can be represented by a sequence  $\mathbf{x} \in [-1, 1]^{d \cdot f_s}$ . An audio neural codec intends to compress  $\mathbf{x}$  and then recover it as  $\hat{\mathbf{x}}$  using an encoder-decoder module with a quantizer:

$$\mathbf{h} = E(\mathbf{x}) \in \mathcal{R}^{T \times L}; \quad \hat{\mathbf{h}} = Q(\mathbf{h}); \quad \hat{\mathbf{x}} = D(\hat{\mathbf{h}}) \quad (1)$$

where  $T$  denotes the number of audio frames after down-sampling in the encoder, and  $L$  denotes the feature dimension of the encoder.  $E$ ,  $D$ , and  $Q$  represent the Encoder, Decoder and Quantizer. Given any frame of encoder output  $\mathbf{h}_t$ , the token sequence  $\mathbf{z}_t = [z_t^1, \dots, z_t^{N_q}]$  is generated by Residual Vector Quantization (RVQ) [14], where  $N_q$  denotes the number of vector quantization layers. With the discrete representation  $\mathbf{z}_t$ ,  $\hat{\mathbf{h}}_t$  is reconstructed as a close estimation of  $\mathbf{h}_t$ , then  $\mathbf{h}_t$  is used to recover  $\mathbf{x}_t$  with the decoder. For an audio sample with  $T$  frame, its discrete representation  $\mathbf{z} \in \mathbf{Z}^{T \times N_q}$  is a matrix. In this work, we convert it into a sequence before being processed by the language model: we use a simple flatten operation along the  $N_q$  dimension. In this work, we set  $N_q = 3$ .

### B. Language Models for Target Audio Separation

The neural audio codec model allows us to operate on discrete audio representations. We train an autoregressive (AR) decoder-only language model for a target audio separation task. The target audio separation task aims to separate the target audio from a mixture of audio based on the query prompt. In general, previous works [8], [13] use a query encoder to extract a global embedding to guide the separation. In this study, we directly use the discrete tokens of query prompt to guide the separation. Specifically, we first use the neural audio codec model to extract the discrete tokens for mixture audio, query prompt, and target audio. Then we combine them as one sequence in order, and apply the language model to it, as Figure 1 (b) shows. Then we can train the language model in two ways. The first is building a prefix-language model

(prefix-LM), which regards the prompt sequence and mixture sequence as the condition to predict the tokens of target audio, it can be formulated as:

$$p(\mathbf{a}|\mathbf{m}, \mathbf{c}; \theta) = \prod_{t=0}^T p(\mathbf{a}_t|\mathbf{a}_{<t}, \mathbf{c}, \mathbf{m}; \theta) \quad (2)$$

where  $\mathbf{a}$  denotes the target audio tokens,  $\mathbf{m}$  and  $\mathbf{c}$  denote the mixture sequence and prompt audio tokens.  $T$  denotes the sequence length of target tokens.  $\theta$  denotes the model parameters. The second way is to build a causal language model (causal-LM) [17], which tries to model the whole sequence in the training process<sup>1</sup>. In the inference, we also regard the prompt and mixture as the condition to predict the target sequence:

$$p(U; \theta) = \prod_{t=0}^N p(\mathbf{u}_t|\mathbf{u}_{<t}; \theta) \quad (3)$$

where  $U = [\mathbf{m}, \mathbf{c}, \mathbf{a}]$ ,  $N$  denotes the total length of the combined sequence. In this work, we adopt the causal-LM for the target audio separation task. We have two considerations: (1) From the training aspect, causal-LM increases the training difficulty and improves the generalization due to it asks LM to model the relationships internally between condition tokens. (2) Our preliminary experimental results show that causal-LM performs better than prefix-LM. In the model structure aspect, we adopt a multi-scale transformer [16] for the language model. To help the language model better understand the query, mixture, and target sequence, we add the special token [a\_start] and [a\_end] encompass the mixture and target sequence, [p\_start] and [p\_end] encompass the prompt sequence.

### C. Pre-training

In the previous section, we introduce how to train a causal language model with the triple data (mixture audio, prompt audio, target audio). However, such a training strategy costs a lot of effort to prepare the large-scale dataset. In this part, we introduce a pre-training strategy to utilize a large-scale audio-only dataset, as Figure 1 (a) shows. Our motivation is that consistency and relevance are the key points for a target audio separation model. *e.g.* The model should be able to output consistent target source audio without any other noise. Furthermore, the model should output the most relevant source audio with the prompt prompt, so the model should learn to understand the most relevant source with the prompt. To improve the consistency and relevance of the language model, we design two pre-training tasks: Audio Continuation and Audio Inpaint. Specifically, for any audio samples, we extract its token sequence using the audio codec model, and then we design two tasks:

**Audio Continuation** We design an audio continuation pre-training task to improve the ability of the language model to generate consistent audio. We let the LM predict the following

<sup>1</sup>Note that causal-LM and prefix-LM have the same model configuration, the difference is attention mask strategy and loss calculation.

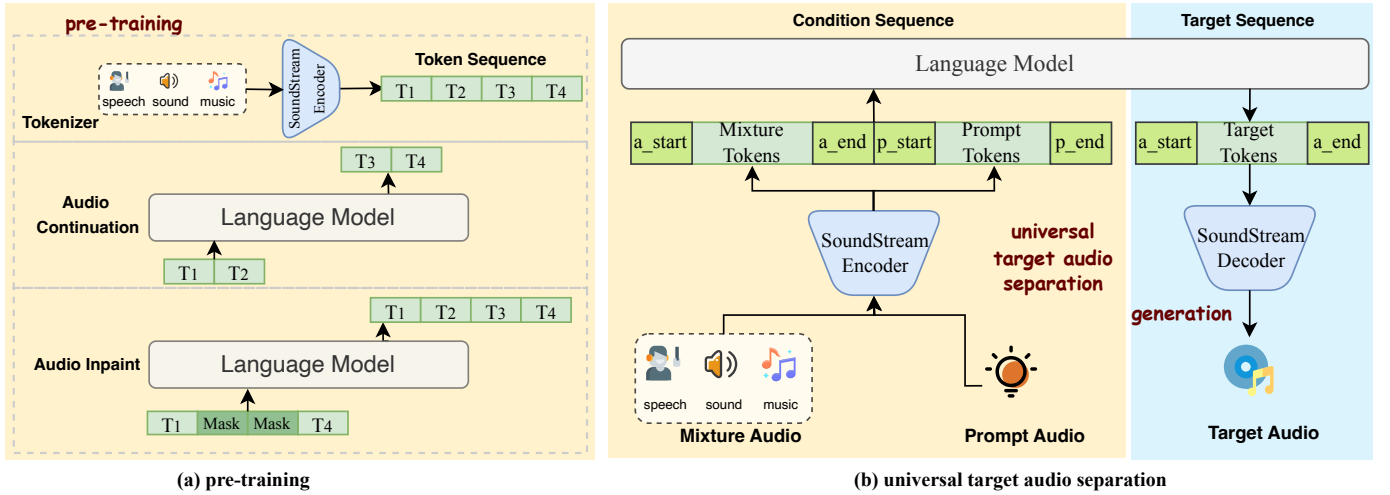


Fig. 1. (a) gives the illustration of two pre-training tasks. (b) shows the Sequence layout for UniSep. We encode all input into the discrete token space so that we can directly use language model architectures for audio separation.

TABLE I  
THE TRAINING DATASET USED IN UNISEP. AUDIO-ONLY DENOTES WE USE THE ORIGINAL AUDIO WITHOUT ANY DATA SIMULATION. SIMULATION DENOTES WE COME MULTIPLE AUDIO AS ONE MIXTURE AUDIO.

Dataset	Source type	Duration(hrs)	
		audio-only	simulation
LibriLight [18]	Speech	20k	-
MLS [20]		-	10k
AudioSet [19]	Sound	5.8k	10.1k
MUSDB18 [21]	Music	-	722

audio sequence based on the previous part.

**Audio Inpaint** We design an audio inpaint pre-training task to improve the ability of the language model to understand the relevance between different audio segments. We add a [MASK] token and randomly replace some tokens in the original sequence (with probability 20% to mask). Then, we use the masked sequences as conditions, to predict the masked token and recover the original sequence.

### III. EXPERIMENTAL SETUP

#### A. Dataset

As shown in Table I, our model uses five publicly available datasets with a total duration of 36.5k hours. We utilize the LibriLight [18] and AudioSet [19] for pre-training and employ the MLS [20], AudioSet [19], and MUSDB18 [21] for training the audio separation task.

**Pre-training.** In our experiments, we utilize the unsupervised LibriLight datasets for pre-training, which are more easily accessible than labeled data. For pre-training, we divide the data into 10-second audio segments, using a total of 20k hours from LibriLight and 5.8k hours from AudioSet. This pre-training takes about 2 weeks.

**Speech.** For our target speech separation task, we randomly mixed two distinct segments from different speakers, with a -5 to 10 SNR. Each mixed and target speech lasts between 3-10 seconds. We also extract 3 seconds from other target speaker segments as prompt speech. In this way, we simulate approximately 10k hours of MLS for the target speech separation

task. We test the models' performance on the test set of the Libri2Mix test set [22], sampled at 16 kHz.

**Sound.** AudioSet is a weakly labeled dataset, which has complex caption annotations of each audio clip but lacks accurate timestamps. For target audio separation, Some researchers [13] pointed out that due to the complexity of audio clips, it is difficult to find a semantically consistent but different audio sample to serve as the query audio. Therefore, we utilize similar methods to simulate data in the AudioSet. We follow the pipeline proposed in [13], which uses a Sound Event Detection (SED) system to detect each audio clip and extract more accurate sound segments. Next, we simulate the dataset using following ways: (1) We create mixtures by randomly selecting two audio clips with different sound events and simply using the target source in the mixture as prompt audio. (2) We use the SED to detect 4-second audio segments, with the first 2 seconds as prompt and the subsequent 2 seconds as target. We simulate approximately 10.1k hours of mixed data, with a signal-to-noise ratio (SNR) ranging from -5 to 10 in both cases.

**Music.** MUSDB18 has four different genres along with their isolated drums, bass, vocals, and others stems. In our approach, we randomly select a 3-second prompt from a specific music track, followed by 5-8 seconds designated as the target. Subsequently, we randomly choose 1-2 stems from the other three categories to mix. We simulate 722 hours as the train set and 1.44 hours as the test set in target music separation.

#### B. System Implementation

**Implementation details.** Our UniSep is trained with AdamW optimizer ( $\beta_1=0.9$ ,  $\beta_2=0.95$ ,  $\text{eps}=1\text{e-}8$ ) with an initial learning rate of  $5\text{e-}4$ . We perform the same learning rate decay strategy with [25] and the warm-up step is 10000. Our batch scale is 4800 tokens, which represents the sum of sequence length. The transformer architecture in our UniSep comprises 12 layers for global transformer and 4 layers for local transformer, with each layer containing 8 attention heads. The embedding dimension for the transformer is set to 1536. Our model has a total of

TABLE II  
PERFORMANCE EVALUATION ON TARGET SPEECH SEPARATION.

Model	Pre-train	Libri2Mix test set(↑)		
		PESQ	DNSMOS	MUSHRA
SpeakerBeam [23]		<b>2.89</b>	3.18	3.68±0.10
VoiceFilter [24]		2.41	3.35	3.43±0.09
UniAudio(single) [16]	N	1.97	3.93	3.58±0.08
UniAudio [16]		1.88	3.96	3.72±0.06
UniSep(Single)	N	1.85	3.89	3.54±0.13
UniSep(Single)	Y	2.09	3.96	3.71±0.14
UniSep	N	2.14	4.01	3.89±0.11
<b>UniSep</b>	<b>Y</b>	<b>2.23</b>	<b>4.07</b>	<b>4.11±0.11</b>

535 million parameters and is trained on 8 Tesla V100 GPUs. We use a pre-trained neural audio codec model<sup>2</sup> to extract discrete tokens, which is based on SoundStream. To eliminate the impact of the codec, we use the audio reconstructed by the codec as the reference audio during evaluation.

**System evaluation.** To measure signal quality, our UniSep is evaluated using both objective and subjective metrics. Perceptual Evaluation of Speech Quality (PESQ) and non-intrusive Deep Noise Suppression Mean Opinion Score (DNS-MOS) [26] are signal-level objective quality metrics derived from human auditory research. ViSQOL (Virtual Speech Quality Objective Listener) [27], [28] is an objective metric for assessing perceived audio quality. It compares the spectro-temporal similarity between a reference and a test speech signal to generate a MOS-LQO (Mean Opinion Score - Listening Quality Objective) score, which ranges from 1 (worst) to 5 (best). STFT distance [29] measures the loss between the log magnitude spectrograms of the reference and ground truth waveforms. We use PESQ and DNSMOS as objective metrics for evaluating speech, and ViSQOL and STFT distance as objective metrics for assessing music and sound. Note that since our UniSep uses autoregressive sampling at their output, metrics such as PESQ and scale-invariant signal-to-noise ratio improvement (SI-SNRi) that rely on exact sample alignment between references and estimates are not suitable [30]. Moreover, we also perform a subjective test to measure the performance of our models, which can provide a more reliable measurement of relative quality. We conduct a Multiple Stimuli with Hidden Reference and Anchor (MUSHRA) listening test which is designed to assess the perceived quality of different audio samples under specific conditions. We use the MUSHRA as the subjective test that primarily evaluates the similarity between generated audio and target audio (ground truth). We randomly selected 50 sentences from the test set and recruited 20 people to evaluate the audio. The score range is from 1 to 5.

#### IV. RESULTS AND ANALYSES

UniSep is first trained on the segmented audio-only LibriLight dataset with about 20k hours. Next-token prediction and mask techniques accomplish audio continuation and audio inpainting, which also enhance the language model’s comprehension capabilities. Based on the pre-trained model,

<sup>2</sup><https://github.com/yangdongchao/UniAudio>

TABLE III  
PERFORMANCE EVALUATION ON TARGET SOUND SEPARATION.

Model	Pre-train	AudioSet test set		
		STFT(↓)	ViSQOL(↑)	MUSHRA(↑)
CLAPSep [31]	N	4.13	2.99	-
SoloAudio [32]	N	2.86	3.33	-
USS-ResUNet30 [33]	N	2.28	3.31	3.2±0.27
UniSep(Single)	N	1.43	3.86	2.95±0.15
UniSep(Single)	Y	1.35	3.94	3.24±0.13
UniSep	N	1.28	4.02	3.55±0.10
<b>UniSep</b>	<b>Y</b>	<b>1.24</b>	<b>4.10</b>	<b>3.73±0.08</b>

we train the target audio separation task on the simulated MLS, AudioSet, ESC-50, and MUSDB18 datasets. Ultimately, we can use a universal model to perform universal audio separation for speech, sound, and music. UniSep (Single) represents the training of each task using our model structure in Table II III IV, so there are three separate single-task models. We can see that single-task models are poor than unified model, which validated the importance of building a unified separation model. In the following, we evaluate our UniSep on speech, sound, and music test set.

##### A. Results on Speech

For speech separation, we use the Libri2Mix test set to evaluate the target speech separation task. As shown in Table II, our UniSep performs better than the single-task model in both objective metrics and subjective evaluations. We can also observe that the model performance is improved with pre-training, which validates the effectiveness of the pre-training methods we proposed. Furthermore, compared with previous works, such as UniAudio, VoiceFilter, and SpeakerBeam, UniSep also shows better performance, *e.g.* better DNSMOS and MUSHRA metrics. We also observe a similar finding with previous work [30]: the objective metric PESQ is may not suitable for a sampling-based generation model, due to it depending on the strict alignment between reference and generated speech.

##### B. Results on Sound

In this section, we evaluate our UniSep model on the AudioSet test set. From the Table III, we can conclude the same results with Section IV-A. UniSep with pre-training significantly outperforms different baselines across multiple metrics, validating UniSep’s superiority. Due to time and resource constraints, we performed only demo-level comparisons (average) on CLAPSep [31] and SoloAudio [32]. Their weaker performance could be due to a lack of specialized training on music datasets.

Additionally, it’s important to mention that we used the AudioSet dataset in this experiment, which includes speech, music, and sound events, with speech and music making up the majority. Therefore, during the training and inference stages, any combination of these three types of data could exist. In the AudioSet section of our demo page, one can listen to



TABLE IV  
PERFORMANCE EVALUATION ON TARGET MUSIC SEPARATION(MUSDB18 TEST SET). AVG REPRESENTS THE AVERAGE EVALUATION RESULTS ACROSS ALL CATEGORIES. STFT MEANS STFT DISTANCE.

Model	Pre-train	STFT(↓)	ViSQOL(↑)					MUSHRA(↑)				
			vocals	drums	bass	other	avg	vocals	drums	bass	other	avg
CLAPSep [31]	N	5.78	-	-	-	-	3.60	-	-	-	-	-
SoloAudio [32]	N	2.90	-	-	-	-	3.58	-	-	-	-	-
UniSep(Single)	N	1.82	3.99	4.03	4.08	4.04	4.04	3.16±0.10	2.84±0.14	2.86±0.14	3.03±0.13	2.97
UniSep(Single)	Y	1.63	4.17	4.06	4.12	4.08	4.11	3.52±0.14	3.03±0.12	3.14±0.18	3.26±0.11	3.24
UniSep	N	1.33	4.22	4.19	4.18	4.21	4.20	3.74±0.09	3.41±0.09	3.33±0.09	3.60±0.07	3.52
<b>UniSep</b>	<b>Y</b>	<b>1.21</b>	<b>4.26</b>	<b>4.22</b>	<b>4.19</b>	<b>4.24</b>	<b>4.23</b>	<b>3.98±0.11</b>	<b>3.55±0.07</b>	<b>3.50±0.08</b>	<b>3.85±0.12</b>	<b>3.72</b>

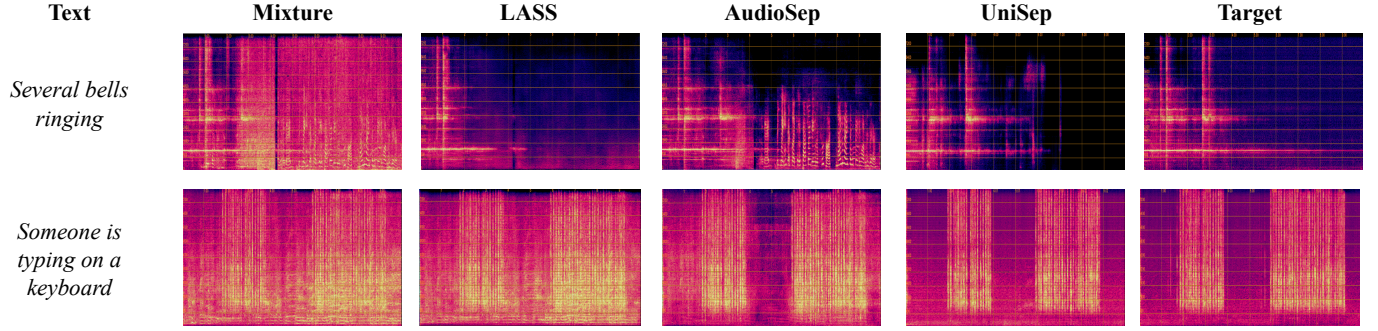


Fig. 2. Visualization results on language-queried audio source separation.

examples of music mixed with speech or sound events, which also demonstrates the universality of our UniSep.

### C. Results on Music

As demonstrated in Table IV, we explore the separation performance of different stems in MUSDB18. Overall, the subjective and objective metrics of the three experiments demonstrate that our method can generate audio with better perceptual quality. The average metrics indicate that the results are essentially consistent with both speech and sound. This suggests that universal multitasking outperforms single-task approaches, and the pre-training method we employed effectively enhances the model’s performance.

### D. UniSep as foundation model for downstream task

UniSep can be fine-tuned using a small amount of data to accomplish other tasks, *e.g.* language-queried audio source separation task [34], [35]. We show that UniSep can be fine-tuned on text-audio pairs, and then it can be used to separate audio based on text description. We utilize T5 [36] to encode the text description. The encoded text is then used as a query to guide the separation process, allowing UniSep to isolate the audio sources that match the given textual description.

Table V shows the results, where our UniSep, after fine-tuning, performs slightly worse compared to AudioSep. It is worth noting that AudioSep is trained on a massive dataset of 14,000 hours of audio-text pairs. In contrast, our UniSep is fine-tuned on a much smaller dataset of only around 16.7 hours from the AudioCaps dataset. Compared to AudioSep, UniSep achieves 96.94% of ViSQOL performance with only 0.12% audio-text pairs. Under identical data conditions, compared to LASS, UniSep achieved a relative performance improvement of 36.09% in STFT, demonstrating its efficient performance

TABLE V  
PERFORMANCE EVALUATION ON LANGUAGE-QUERIED AUDIO SOURCE SEPARATION. ‘HOURS’ REFERS TO THE VOLUME OF THE AUDIO-TEXT PAIRS DATASET USED DURING THE TRAINING PROCESS. UNISEP HAS THE SAME AMOUNT OF AUDIO-TEXT DATA AS LASS.

Model	Hours	AudioCaps test set	
		STFT(↓)	ViSQOL(↑)
AudioSep [35]	14k	<b>1.97</b>	<b>3.60</b>
LASS [34]	17.3	3.63	3.40
UniSep+Fine-tuning	16.7	2.32	3.49

for downstream tasks. We also visualize the spectrograms for audio mixtures, target audio sources, and separated sources using different models, as depicted in Figure 2. The spectrogram patterns of the separated sources resemble those of the target sources, which highlights the model’s effectiveness. Audio samples can be found on our project page.

## V. CONCLUSION

We propose a universal target audio separation model, UniSep, based on language models. Different from previous separation models that try to separate the audio in continuous space *e.g.* spectrogram, we propose to separate the audio in a discrete latent space. Then we propose to use language model for target audio separation task. We also propose a novel pre-training strategy to use large-scale audio-only datasets, which significantly reduces the efforts of data simulation. To the best of our knowledge, UniSep is the first unified model that can separate any audio source without considering the audio types and the number of audio sources. Extensive experiments show the effectiveness of UniSep, we demonstrate the effectiveness of scaling data in audio separation. In the future, we can improve UniSep from two aspects: (1) Training a better audio tokenizer; (2) Scaling the model size and data quantity.

## REFERENCES

- [1] Aditya Arie Nugraha, Antoine Liutkus, and Emmanuel Vincent, "Multi-channel audio source separation with deep neural networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 9, pp. 1652–1664, 2016.
- [2] Prem Seetharaman, Gordon Wichern, Shrikant Venkataramani, and Jonathan Le Roux, "Class-conditional embeddings for music source separation," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 301–305.
- [3] DeLiang Wang and Jitong Chen, "Supervised speech separation based on deep learning: An overview," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 10, pp. 1702–1726, 2018.
- [4] Yi Luo and Nima Mesgarani, "Conv-tasnet: Surpassing ideal time-frequency magnitude masking for speech separation," *IEEE/ACM transactions on audio, speech, and language processing*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [5] Ilya Kavalerov, Scott Wisdom, Hakan Erdogan, Brian Patton, Kevin Wilson, Jonathan Le Roux, and John R. Hershey, "Universal sound separation," in *2019 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2019, pp. 175–179.
- [6] Efthymios Tzinis, Scott Wisdom, John R Hershey, Aren Jansen, and Daniel PW Ellis, "Improving universal sound separation using sound classification," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 96–100.
- [7] Scott Wisdom, Efthymios Tzinis, Hakan Erdogan, Ron Weiss, Kevin Wilson, and John Hershey, "Unsupervised sound separation using mixture invariant training," *Advances in Neural Information Processing Systems*, vol. 33, pp. 3846–3857, 2020.
- [8] Yuanyuan Wang, Hangting Chen, Dongchao Yang, Jianwei Yu, Chao Weng, Zhiyong Wu, and Helen Meng, "Consistent and relevant: Rethink the query embedding in general sound separation," in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 961–965.
- [9] Alexandre Défossez, Nicolas Usunier, Léon Bottou, and Francis Bach, "Music source separation in the waveform domain," *arXiv preprint arXiv:1911.13254*, 2019.
- [10] Simon Rouard, Francisco Massa, and Alexandre Défossez, "Hybrid transformers for music source separation," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [11] Jordi Pons, Xiaoyu Liu, Santiago Pascual, and Joan Serra, "Gass: Generalizing audio source separation with large-scale data," *arXiv preprint arXiv:2310.00140*, 2023.
- [12] Scott Wisdom, Hakan Erdogan, Daniel PW Ellis, Romain Serizel, Nicolas Turpault, Eduardo Fonseca, Justin Salamon, Prem Seetharaman, and John R Hershey, "What's all the fuss about free universal sound separation data?," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 186–190.
- [13] Ke Chen, Xingjian Du, Bilei Zhu, Zejun Ma, Taylor Berg-Kirkpatrick, and Shlomo Dubnov, "Zero-shot audio source separation through query-based learning from weakly-labeled data," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022, vol. 36, pp. 4441–4449.
- [14] Neil Zeghidour, Alejandro Luebs, Ahmed Omran, Jan Skoglund, and Marco Tagliasacchi, "Soundstream: An end-to-end neural audio codec," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 495–507, 2021.
- [15] Zalán Borsos, Raphaël Marinier, Damien Vincent, Eugene Kharitonov, Olivier Pietquin, Matt Sharif, Dominik Roblek, Olivier Teboul, David Grangier, Marco Tagliasacchi, et al., "Audiolm: a language modeling approach to audio generation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2023.
- [16] Dongchao Yang, Jinchuan Tian, Xu Tan, Rongjie Huang, Songxiang Liu, Xuankai Chang, Jiatong Shi, Sheng Zhao, Jiang Bian, Xixin Wu, et al., "Uniaudio: An audio foundation model toward universal audio generation," *arXiv preprint arXiv:2310.00704*, 2023.
- [17] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al., "Improving language understanding by generative pre-training," 2018.
- [18] Jacob Kahn, Morgane Riviere, Weiyi Zheng, Evgeny Kharitonov, Qiantong Xu, Pierre-Emmanuel Mazaré, Julien Karadayi, Vitaliy Liptchinsky, Ronan Collobert, Christian Fuegen, et al., "Libri-light: A benchmark for asr with limited or no supervision," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7669–7673.
- [19] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter, "Audio set: An ontology and human-labeled dataset for audio events," in *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 776–780.
- [20] Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, and Ronan Collobert, "Mls: A large-scale multilingual dataset for speech research," *arXiv preprint arXiv:2012.03411*, 2020.
- [21] Zafar Rafii, Antoine Liutkus, Fabian-Robert Stöter, Stylianos Ioannis Mimilakis, and Rachel Bittner, "Musdb18-a corpus for music separation," 2017.
- [22] Joris Cosentino, Manuel Pariente, Samuele Cornell, Antoine Deleforge, and Emmanuel Vincent, "Librimix: An open-source dataset for generalizable speech separation," *arXiv preprint arXiv:2005.11262*, 2020.
- [23] Kateřina Zmolíková, Marc Delcroix, Keisuke Kinoshita, Tsubasa Ochiai, Tomohiro Nakatani, Lukáš Burget, and Jan Černocký, "Speakerbeam: Speaker aware neural network for target speaker extraction in speech mixtures," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 4, pp. 800–814, 2019.
- [24] Quan Wang, Hannah Muckenhirn, Kevin Wilson, Prashant Sridhar, Zelin Wu, John Hershey, Rif A Saurous, Ron J Weiss, Ye Jia, and Ignacio Lopez Moreno, "Voicefilter: Targeted voice separation by speaker-conditioned spectrogram masking," *arXiv preprint arXiv:1810.04826*, 2018.
- [25] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [26] Chandan KA Reddy, Vishak Gopal, and Ross Cutler, "Dnsmos: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6493–6497.
- [27] Andrew Hines, Jan Skoglund, Anil C Kokaram, and Naomi Harte, "Visqol: an objective speech quality model," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2015, pp. 1–18, 2015.
- [28] Michael Chinen, Felicia SC Lim, Jan Skoglund, Nikita Gureev, Feargus O'Gorman, and Andrew Hines, "Visqol v3: An open source production ready objective speech and audio metric," in *2020 twelfth international conference on quality of multimedia experience (QoMEX)*. IEEE, 2020, pp. 1–6.
- [29] Rithesh Kumar, Prem Seetharaman, Alejandro Luebs, Ishaan Kumar, and Kundan Kumar, "High-fidelity audio compression with improved rvqgan," 2023.
- [30] Hakan Erdogan, Scott Wisdom, Xuankai Chang, Zalán Borsos, Marco Tagliasacchi, Neil Zeghidour, and John R Hershey, "Tokensplit: Using discrete speech representations for direct, refined, and transcript-conditioned speech separation and recognition," *arXiv preprint arXiv:2308.10415*, 2023.
- [31] Hao Ma, Zhiyuan Peng, Xu Li, Mingjie Shao, Xixin Wu, and Ju Liu, "Clapsep: Leveraging contrastive pre-trained model for multi-modal query-conditioned target sound extraction," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2024.
- [32] Helin Wang, Jiarui Hai, Yen-Ju Lu, Karan Thakkar, Mounya Elhilali, and Najim Dehak, "Soloaudio: Target sound extraction with language-oriented audio diffusion transformer," *arXiv*, 2024.
- [33] Qiuqiang Kong, Ke Chen, Haohe Liu, Xingjian Du, Taylor Berg-Kirkpatrick, Shlomo Dubnov, and Mark D Plumbley, "Universal source separation with weakly labelled data," *arXiv preprint arXiv:2305.07447*, 2023.
- [34] Xubo Liu, Haohe Liu, Qiuqiang Kong, Xinhao Mei, Jinzheng Zhao, Qiushi Huang, Mark D Plumbley, and Wenwu Wang, "Separate what you describe: Language-queried audio source separation," *arXiv preprint arXiv:2203.15147*, 2022.
- [35] Xubo Liu, Qiuqiang Kong, Yan Zhao, Haohe Liu, Yi Yuan, Yuzhuo Liu, Rui Xia, Yuxuan Wang, Mark D Plumbley, and Wenwu Wang, "Separate anything you describe," *arXiv preprint arXiv:2308.05037*, 2023.
- [36] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu, "Exploring the limits of transfer learning with a unified text-to-text transformer," *Journal of Machine Learning Research*, vol. 21, no. 140, pp. 1–67, 2020.