

# Leveraging Language Information for Target Language Extraction

Mehmet Sinan Yıldırım\*, Ruijie Tao\*, Wupeng Wang\*, Junyi Ao<sup>†</sup>, Haizhou Li\*<sup>†</sup>

\* Department of Electrical and Computer Engineering, National University of Singapore, Singapore

E-mail: sinan@u.nus.edu

<sup>†</sup> School of Artificial Intelligence, Shenzhen Research Institute of Big Data,  
The Chinese University of Hong Kong, Shenzhen, China

**Abstract**—Target Language Extraction aims to extract speech in a specific language from a mixture waveform that contains multiple speakers speaking different languages. The human auditory system is adept at performing this task with the knowledge of the particular language. However, the performance of the conventional extraction systems is limited by the lack of this prior knowledge. Speech pre-trained models, which capture rich linguistic and phonetic representations from large-scale in-the-wild corpora, can provide this missing language knowledge to these systems. In this work, we propose a novel end-to-end framework to leverage language knowledge from speech pre-trained models. This knowledge is used to guide the extraction model to better capture the target language characteristics, thereby improving extraction quality. To demonstrate the effectiveness of our proposed approach, we construct the first publicly available multilingual dataset for Target Language Extraction. Experimental results show that our method achieves improvements of 1.22 dB and 1.12 dB in SI-SNR for English and German extraction, respectively, from mixtures containing both languages.

**Index Terms**—Target Language Extraction, Multilingual Speech Separation, Pre-trained Speech Models

## I. INTRODUCTION

Humans have the extraordinary ability to focus on speech from a particular auditory source in a noisy environment with various interfering sources. This ability is known as the cocktail-party effect [1]. Inspired by this phenomenon, enabling machines to replicate this selective auditory perception has long been a significant research challenge. Recent advances in deep learning have led to a paradigm shift, enabling the development of high-quality models for this task.

Target speech extraction models are a common solution designed to mimic this selective auditory perception capability. These models are generally conditioned on an additional input from which they obtain the speaker characteristics. This auxiliary input is usually a reference utterance produced by the target speaker [2]–[5], but it can alternatively take the form of textual cues [6]–[9], visual cues [10]–[12], or even EEG signals [13].

The aforementioned cues are typically used in monolingual settings. However, real-world environments often involve

This research is supported by: Shenzhen Science and Technology Program (Shenzhen Key Laboratory, Grant No. ZDSYS20230626091302006); Shenzhen Science and Technology Research Fund (Fundamental Research Key Project, Grant No. JCYJ20220818103001002); Program for Guangdong Introducing Innovative and Entrepreneurial Teams, Grant No. 2023ZT10X044.




Fig. 1. Overview of the Target Language Extraction (TLE) task: given a speech mixture containing multiple languages (e.g. English and German), the model extracts only the speech corresponding to the target language (e.g., English), regardless of speaker identity.

multiple languages. For example, airport announcements are often delivered in several languages, and international participants at academic conferences frequently interact in different languages. In such scenarios, the user may be interested in extracting speech from a specific language. To address this need, the Target Language Extraction (TLE) task (as shown in Fig. 1) was introduced [14] with the aim of isolating speech corresponding to a particular language from multilingual conversations. This task can serve as an essential preprocessing step in multilingual settings, enabling downstream systems to process each language separately.

To achieve TLE, prior works [14]–[16] primarily follow approaches developed for Target Speaker Extraction (TSE), adapting these models to extract languages rather than speakers. Representative models include SpEx+ [17], Conv-TasNet [18], and SepFormer [19]. However, these methods largely ignore the inherent linguistic structure of the target language. Studies in auditory perception suggest that humans recognize and extract speech more effectively in familiar languages than in unfamiliar ones [20], due to prior exposure to the language’s phonological and syntactic patterns. Natural languages exhibit unique phonetic and linguistic characteristics. Familiarity with these characteristics can aid with extraction and help mitigate potential errors during the extraction process.

Motivated by the role of linguistic knowledge in human auditory perception, we hypothesize that injecting language information into the extraction network can enhance the overall quality of TLE. A key challenge is how to incorporate such language information effectively into the extraction process. Re-

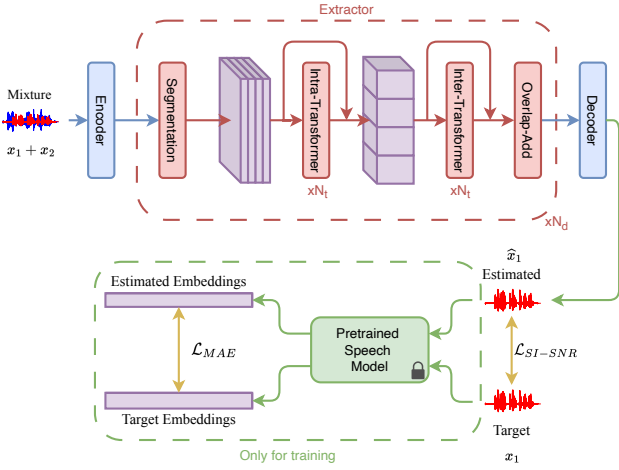


Fig. 2. Overview of the proposed method. Pre-trained speech model obtains embeddings for both the estimated and the target waveforms. An additional similarity loss between them contributes to align the internal speech representations of the base model

cent advances in self-supervised speech representation learning provide a promising solution, as these models are trained on large-scale unlabeled data and are capable of capturing universal, robust speech and language representations [21]–[23]. Large-scale pre-trained speech models have demonstrated a strong capacity for language modeling, capturing not only phonetic features but also higher-level linguistic structures. To facilitate TLE with language information, we introduce a novel approach that employs an auxiliary training loss. This loss aligns the internal representations of the extraction model with those from a pre-trained speech model. No additional computation is incurred during inference since this auxiliary supervision is only applied during training.

In addition, previous TLE studies [14]–[16] constructed their mixture datasets using samples from the GlobalPhone corpus [24], which is not publicly available. To enable a fair comparison on an open-source benchmark, we simulate a new bilingual dataset following the same principles using the CommonVoice [25] dataset which is publicly available.

In summary, our contributions are three-fold:

- We propose a framework that injects language information into the TLE pipeline by leveraging pre-trained speech models.
- We provide the first publicly available dataset for TLE, named CommonVoiceMix
- We demonstrate that our framework enhances the extraction performance for TLE and establish it as the first publicly available baseline for future studies

The remainder of this manuscript is organized as follows. In section II, we present a formal definition of TLE, the design of the combined loss function, the newly constructed dataset, and the experimental setup. Section III presents the experimental results and their interpretations. Section IV concludes this work.

## II. TARGET LANGUAGE EXTRACTION WITH LANGUAGE INFORMATION

### A. Target Language Extraction

Target Language Extraction is the task of isolating speech in a specific target language from a mixture containing speech in multiple languages by different speakers. Let  $x_i(t)$  denote the speech waveform for language  $i$ . Then the mixture signal  $m(t)$  containing speech from  $N$  languages is given by:

$$m(t) = \sum_{i=1}^N x_i(t)$$

TLE aims to recover the original target language speech from this mixture. This can be formulated as finding an optimal transformation  $f(\cdot)$  such that the estimation

$$\hat{x}_i(t) = f(m(t))$$

is as close to the original target language speech  $x_i(t)$  as possible according to a certain distance metric.

Previous works proposed using the scale-invariant signal-to-noise ratio (SI-SNR) [26]. We adopt this as our baseline loss. If the target language is language  $i$ , the SI-SNR loss is defined as

$$\mathcal{L}_{SI-SNR} = -10 \log_{10} \left( \frac{\|\alpha x_i(t)\|^2}{\|\alpha x_i(t) - \hat{x}_i(t)\|^2} \right)$$

where

$$\alpha = \frac{\hat{x}_i(t)^T x_i(t)}{\|x_i(t)\|^2}.$$

The SI-SNR loss ensures that the estimated output closely matches the target output in terms of signal fidelity.

### B. Proposed Method for Injecting Language Information

This vanilla TLE approach does not explicitly use any intrinsic information of the underlying natural language, as it only optimizes for fidelity, although leveraging such information could benefit the extraction process. Self-supervised speech models capture this information as they are trained on large speech corpora to learn latent representations for speech signals. Since we work with multiple languages, it is more beneficial to use models trained on multilingual corpora that include the target languages in our dataset.

To achieve this goal, we introduce a TLE model that incorporates language knowledge as supervision to better guide the extraction process. Our model is inspired by the Sepformer [19], which has shown strong performance on monaural speech processing. As shown in Fig. 2, the mixture speech containing multiple languages is processed with a 1-D convolutional encoder. Then the encoder output is segmented along the time dimension into chunks of fixed length. An Intra-Transformer module applies attention within each chunk to model local dependencies. Subsequently, an Inter-Transformer module applies attention across corresponding time points of different chunks to capture global dependencies. The segments are then

recombined and decoded to obtain the target speech estimate  $\hat{x}_i(t)$ .

We process both the target signal  $x_i(t)$  and the estimated signal  $\hat{x}_i(t)$  through a self-supervised pre-trained speech model to obtain their latent representations. These representations are rich in contextual information learned from the dependencies in the pre-training corpus. Ideally, these two representations should be as similar as possible. To enforce this similarity, we define a distance function between them as an auxiliary loss to guide the model toward extracting speech with correct linguistic structure and thus achieving better language understanding. We use the L1 distance to this end. Since the SI-SNR loss is in logarithmic scale, we also take the logarithm of the L1 distance to match the scale and sensitivity of the losses. Therefore, given a certain pre-trained model  $h(\cdot)$ , we have:

$$\mathcal{L}_{MAE} = 10 \log_{10} \left( \frac{1}{d} \sum_{j=1}^d \left| [h(x_i(t))]_j - [h(\hat{x}_i(t))]_j \right| \right)$$

where  $d$  is the dimension of the pre-trained model output. We add this loss to the original SI-SNR loss with a coefficient  $\beta \geq 0$ .

$$\mathcal{L} = \mathcal{L}_{SI-SNR} + \beta \mathcal{L}_{MAE}$$

$\beta = 0$  corresponds to the baseline case of only using SI-SNR loss. The overall architecture is illustrated in Fig. 2. We train the model in two stages: first with  $\beta = 0$  to learn general extraction, and then with a positive  $\beta$  value to refine the extraction using language information. Note that this auxiliary supervision does not incur additional computational cost at inference time since it is only used during training.

### C. CommonVoiceMix: A New Public Corpus for TLE

To evaluate our method, we require a suitable multilingual dataset. In the TSE literature, commonly used datasets such as LibriMix [27] and wsj0-2mix [28] generate mixture signals from the LibriSpeech [29] and WSJ0 [30] corpora, respectively. We adopt the LibriMix strategy to construct our dataset.

To construct multilingual mixtures, we need a base dataset that contains utterances in different languages and is ideally publicly available. Additionally, it should contain a high number of speakers to ensure that the method focuses on language rather than speaker identity. Due to these reasons, we chose the publicly available CommonVoice dataset [25]. The appropriate mixtures are obtained by mixing samples from different languages. We chose English and German for our experiments. We discarded samples shorter than 7 seconds. The speaker IDs are utilized to verify that there is no speaker overlap between the train, dev, and test sets of the two languages.

With this strategy, we constructed the CommonVoiceMix dataset that comprises 30000, 4600 and 4500 English-German mixtures for train, dev, and test sets, respectively. We limit the number of training samples to 30000 whereas dev and test sets are limited by the number of samples of the lower-resource language in the original dataset. We randomly sample 6-second segments from the train mixtures such that the target speech

TABLE I  
A SUMMARY OF THE COMMONVOICEMIX DATASET

Subset	train	dev	test
Total length (h)	50	11.2	11.2
# of samples	30000	4600	4500
# of unique English speakers	9945	3126	4000
# of unique German speakers	4203	2276	2739

is not empty. The entire utterance is used for dev and test sets. This corresponds to approximately 50 hours, 11.2 hours, and 11.2 hours for train, dev, and test sets, respectively. We also utilize the client\_id values of the speakers to identify the number of unique speakers in the dataset. A summary can be seen in Table I.

We believe that a public TLE dataset will promote and accelerate research in this field. By using the same dataset, new approaches can directly compare to our work, which can serve as a baseline. For this purpose, our code and metadata are publicly available.<sup>1</sup> Other researchers can use these to construct our dataset and also create mixtures of other language combinations available in the CommonVoice dataset.

### D. Experimental Setup

[14] shows that the best baseline for TLE is a Sepformer-based model [19]. We use a dual-path transformer-based Sepformer model implemented in Speechbrain [31] using a single mask to generate a single output. We use the default parameter values defined in Speechbrain<sup>2</sup>. However, we set the number of dual computing blocks to 1 due to computational constraints. This corresponds to  $N_t = 8$  and  $N_d = 1$  in Fig. 2. Additionally, the batch size is 2 during training in our experiments. This model serves as the extraction pipeline in our work.

For the self-supervised pre-trained model, there are several options such as HuBERT [21], WavLM [22], and wav2vec 2.0 [23]. As we deal with multiple languages, we choose mHuBERT-147 [32], which is a HuBERT-based self-supervised learning model trained on 147 languages, including both English and German. mHuBERT-147 has the same size as the base HuBERT model. We use the output of the last layer, i.e. the 12th layer.

We use an Adam optimizer with an initial learning rate of  $3 \cdot 10^{-4}$ . We reduce the learning rate by half if the loss has not decreased for three consecutive epochs. We stop the training if the loss has not decreased for 20 epochs in a row. Initially, we start with  $\beta = 0$ . This usually converges in around 80 epochs, which serves as the baseline. Subsequently, we reinitialize the optimizer, set  $\beta = 1$ , and restart training. This second stage takes around 40 epochs to converge.

<sup>1</sup><https://github.com/msinanyildirim/CommonVoiceMix>

<sup>2</sup><https://github.com/speechbrain/speechbrain>

TABLE II  
MAIN RESULTS ON THE COMMONVOICEMIX DEV AND TEST SETS WITH mHUBERT-147 AS THE SPEECH MODEL.  $\uparrow$  MEANS HIGHER IS BETTER.

Set	Target Lang	Method	SI-SNR (dB) $\uparrow$	STOI $\uparrow$	PESQ $\uparrow$
dev	English	baseline	10.85	0.84	1.87
		<b>ours</b>	<b>12.22</b>	<b>0.87</b>	<b>2.08</b>
	German	baseline	10.85	0.87	1.92
		<b>ours</b>	<b>11.98</b>	<b>0.89</b>	<b>2.07</b>
test	English	baseline	9.96	0.82	1.85
		<b>ours</b>	<b>11.18</b>	<b>0.84</b>	<b>2.05</b>
	German	baseline	10.03	0.84	1.91
		<b>ours</b>	<b>11.15</b>	<b>0.87</b>	<b>2.05</b>

### III. EXPERIMENTS AND RESULTS

#### A. Main results

The main experiments are performed as described, with English and German as the target language, respectively. We evaluate the performance in terms of SI-SNR, PESQ [33] and STOI [34], which respectively assess the fidelity, the perceptual quality, and the short-term intelligibility of the estimated signals. For all three metrics, a higher value means a better quality. Table II summarizes the results for English and German. For both languages, the second stage of multitask learning with the mHuBERT-147 significantly improves the results. The SI-SNR results are increased by 1.22 dB and 1.12 dB for English and German extractions, respectively. Similarly, STOI increases by 0.2 and 0.14, while PESQ increases by 0.05 and 0.03, respectively.

These results support our hypothesis. In the second stage, the model optimizes a combined loss that includes the additional language-aware term. This is achieved partly by decreasing the  $\mathcal{L}_{MAE}$ . This occurs when the estimated signals produce latent embeddings similar to those of the target. Because the SSL model embeds language information, the estimated waveforms must be similar to the target waveform in terms of linguistic content and overall contextual consistency. Consequently, the extraction performance improves, as reflected in the higher SI-SNR values and higher PESQ values indicating enhanced perceptual quality.

It is important to note that the SSL model is frozen and not trained together with the core model. Additionally, the core model operations do not depend on the SSL parameters, as the SSL model only processes the output of the core model to calculate the auxiliary loss. Therefore, the SSL model is not necessary during inference. Thus, this performance upgrade with the better language understanding does not affect the inference speed in practical scenarios.

In addition, the overall second stage training does not depend on the specific SSL model used. Different SSL models could be used instead. We compare two models in the next section.

TABLE III  
RESULTS ON COMMONVOICEMIX TEST SET WITH HUBERT AS THE PRE-TRAINED SPEECH MODEL.  $\uparrow$  MEANS HIGHER IS BETTER.

Target Lang	Method	SI-SNR (dB) $\uparrow$	STOI $\uparrow$	PESQ $\uparrow$
English	baseline	9.96	0.82	1.85
	<b>mHuBERT-147</b>	<b>11.18</b>	<b>0.84</b>	<b>2.05</b>
	HuBERT	11.09	0.84	1.97
German	baseline	10.03	0.84	1.91
	<b>mHuBERT-147</b>	<b>11.15</b>	<b>0.87</b>	<b>2.05</b>
	HuBERT	10.97	0.86	1.98

TABLE IV  
RESULTS ON COMMONVOICEMIX TEST SET WITH mHUBERT-147 AS THE SPEECH MODEL FOR DIFFERENT  $\beta$  VALUES.  $\uparrow$  MEANS HIGHER IS BETTER.

$\beta$	English			German		
	SI-SNR (dB) $\uparrow$	STOI $\uparrow$	PESQ $\uparrow$	SI-SNR (dB) $\uparrow$	STOI $\uparrow$	PESQ $\uparrow$
0.1	11.04	0.84	2.01	10.89	0.86	2.01
0.2	10.97	0.84	2.01	11.15	0.86	2.04
<b>1</b>	<b>11.18</b>	<b>0.84</b>	<b>2.05</b>	<b>11.15</b>	<b>0.87</b>	<b>2.05</b>
5	10.83	0.84	2.00	10.83	0.87	2.03
10	10.71	0.84	2.00	10.95	0.87	2.05

#### B. Ablation studies

1) *Using HuBERT*: Instead of mHuBERT-147, we can use the original HuBERT model, which was trained on a corpus containing English speech only. We keep all other parameter values the same. The results are presented in Table III. The SI-SNR improvement over the baseline is comparable to that achieved with the mHuBERT-147 for English. However, for German, the performance gain with HuBERT is slightly less than that with the mHuBERT-147. Nevertheless, there is still a significant performance gain compared to the baseline. This suggests that HuBERT still has fairly good representations for German speech, which helps the extraction. An SSL model trained on speech from both languages is still better. This can be likened to two people: one who only speaks English and another who speaks multiple languages, including English and German. The person who only speaks English can still focus on German speech to some extent, but it is not as effective as the multilingual person.

2) *Different  $\beta$  values*: For our training strategy,  $\beta$  is a hyperparameter that must be chosen carefully to achieve optimal performance. In our results presented in Table IV, we observe that  $\beta \approx 1$  gives a good performance overall. As  $\beta$  deviates from this value, the performance degrades. However, this second stage training strategy still offers improvements over the baseline.

3) *Smaller model size*: Since our strategy injects language information into the extraction process, it can help boost the extraction quality of smaller models. To test this, we try the same procedure for a smaller version of the model with  $N_t = 4$ . This approximately halves the model size. The results are presented in Table V. As expected, the smaller model exhibits lower performance due to its reduced parameter count. In this setting, our training strategy still benefits the extraction

TABLE V

RESULTS ON COMMONVOICEMIX TEST SET USING THE MODEL WITH HALF THE SIZE,  $N_t = 4$  INSTEAD OF 8.  $\uparrow$  MEANS HIGHER IS BETTER.

Target Lang	Method	SI-SNR (dB) $\uparrow$	STOI $\uparrow$	PESQ $\uparrow$
English	baseline	8.17	0.79	1.68
	<b>ours</b>	<b>9.10</b>	<b>0.81</b>	<b>1.83</b>
German	baseline	8.76	0.83	1.74
	<b>ours</b>	<b>9.66</b>	<b>0.85</b>	<b>1.89</b>

TABLE VI

RESULTS ON COMMONVOICEMIX TEST SET USING WAVLM AS THE PRE-TRAINED MODEL.  $\uparrow$  MEANS HIGHER IS BETTER.

Target Lang	Method	SI-SNR (dB) $\uparrow$	STOI $\uparrow$	PESQ $\uparrow$
English	HuBERT	11.09	0.84	1.97
	WavLM	11.13	0.84	2.00
German	HuBERT	10.97	0.86	1.98
	WavLM	11.28	0.87	2.04

performance. However, the absolute SI-SNR gain is lower than in the larger model. This is attributed to the limited capacity of the smaller model, which restricts its ability to process rich representations.

4) *Different pre-trained model:* In our approach, different pre-trained models can be used in place of HuBERT, and each model may result in different extraction quality. As an example, we use WavLM [22] as the pre-trained speech model. A multilingual version of WavLM is not yet publicly available. Therefore, to ensure a fair comparison, we evaluate its performance against the monolingual HuBERT. Additionally, we use WavLM-base since it has a similar number of parameters to HuBERT. As presented in Table VI, we observe that WavLM has a better performance than HuBERT. Our approach is adaptable to future models that may offer even better performance. Finally, a multilingual version of WavLM could potentially perform better, as suggested by the trends in Table III. However, such a model is not currently available for evaluation.

#### IV. CONCLUSION

In this work, we addressed the challenge of Target Language Extraction (TLE), which aims to isolate speech in a target language from multilingual mixtures. We introduced a novel training framework that injects language knowledge into the extraction pipeline by leveraging pre-trained self-supervised speech models. Specifically, we incorporated an auxiliary loss based on mean absolute error between latent representations derived from mHuBERT-147, guiding the model to better capture language-specific characteristics. We validated our approach on CommonVoiceMix, a new publicly available dataset we constructed for this task. Our experiments demonstrated that the proposed method consistently improves extraction quality, achieving notable SI-SNR and perceptual gains, while introducing no additional inference cost. Furthermore, we showed that pre-training on multilingual data enhances performance, highlighting the importance of language diversity in

model design. In the future, we plan to investigate alternative pre-trained models, such as wav2vec 2.0, and extend this framework to causal, real-time scenarios. We believe that our approach provides a promising direction for advancing language-aware speech extraction systems.

#### REFERENCES

- [1] E. C. Cherry, "Some experiments on the recognition of speech, with one and with two ears," *Journal of the acoustical society of America*, vol. 25, pp. 975–979, 1953.
- [2] K. Žmolíková, M. Delcroix, K. Kinoshita, *et al.*, "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.
- [3] Q. Wang, H. Muckenhirn, K. Wilson, *et al.*, "Voicefilter: Targeted voice separation by speaker-conditioned spectrogram masking," in *Interspeech 2019*, 2019, pp. 2728–2732.
- [4] B. Zeng, S. Hongbin, Y. Wan, and M. Li, "Sef-net: Speaker embedding free target speaker extraction network," in *Interspeech 2023*, 2023, pp. 3452–3456.
- [5] H. Sato, T. Moriya, M. Mimura, *et al.*, "Speakerbeam-ss: Real-time target speaker extraction with lightweight conv-tasnet and state space modeling," in *Interspeech 2024*, 2024, pp. 5033–5037.
- [6] M. Huo, A. Jain, C. P. Huynh, *et al.* "Beyond Speaker Identity: Text Guided Target Speech Extraction." arXiv: 2501.09169 [eess]. (Jan. 15, 2025), [Online]. Available: <http://arxiv.org/abs/2501.09169> (visited on 02/19/2025), pre-published.
- [7] X. Liu, Q. Kong, Y. Zhao, *et al.*, "Separate anything you describe," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 33, pp. 458–471, 2025.
- [8] M. Kim, R. Mira, H. Chen, S. Petridis, and M. Pantic, "Contextual speech extraction: Leveraging textual history as an implicit cue for target speech extraction," in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2025, pp. 1–5.
- [9] M. Huo, A. Jain, C. P. Huynh, *et al.*, "Beyond speaker identity: Text guided target speech extraction," in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025, pp. 1–5.
- [10] R. Tao, X. Qian, Y. Jiang, J. Li, J. Wang, and H. Li, "Audio-visual target speaker extraction with selective auditory attention," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 33, pp. 797–811, 2025.
- [11] Z. Mu and X. Yang, "Separate in the speech chain: Cross-modal conditional audio-visual target speech extraction," in *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, ser. IJ-CAI '24, Jeju, Korea, 2024.

- [12] S.-W. Chung, S. Choe, J. S. Chung, and H.-G. Kang, "Facefilter: Audio-visual speech separation using still images," in *Interspeech 2020*, 2020, pp. 3481–3485.
- [13] D. De Silva, S. Cai, S. Pahuja, T. Schultz, and H. Li, "Neurospex: Neuro-guided speaker extraction with cross-modal fusion," in *2024 IEEE Spoken Language Technology Workshop (SLT)*, 2024, pp. 341–348.
- [14] M. Borsdorf, H. Li, and T. Schultz, "Target Language Extraction at Multilingual Cocktail Parties," in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, Cartagena, Colombia: IEEE, Dec. 13, 2021, pp. 717–724.
- [15] M. Borsdorf, K. Scheck, H. Li, and T. Schultz, "Experts Versus All-Rounders: Target Language Extraction for Multiple Target Languages," in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Singapore, Singapore: IEEE, May 23, 2022, pp. 846–850.
- [16] M. Borsdorf, K. Scheck, H. Li, and T. Schultz, "Blind Language Separation: Disentangling Multilingual Cocktail Party Voices by Language," in *Interspeech 2022*, ISCA, Sep. 18, 2022, pp. 256–260.
- [17] M. Ge, C. Xu, L. Wang, E. S. Chng, J. Dang, and H. Li, "Spex+: A complete time domain speaker extraction network," in *Interspeech 2020*, 2020, pp. 1406–1410.
- [18] Y. Luo and N. Mesgarani, "Conv-tasnet: Surpassing ideal time–frequency magnitude masking for speech separation," *IEEE/ACM transactions on audio, speech, and language processing*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [19] C. Subakan, M. Ravanelli, S. Cornell, M. Bronzi, and J. Zhong, "Attention is all you need in speech separation," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2021, pp. 21–25.
- [20] A. R. Bradlow and D. B. Pisoni, "Recognition of spoken words by native and non-native listeners: Talker-, listener-, and item-related factors," *The Journal of the Acoustical Society of America*, vol. 106, no. 4, pp. 2074–2085, 1999.
- [21] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM transactions on audio, speech, and language processing*, vol. 29, pp. 3451–3460, 2021.
- [22] S. Chen, C. Wang, Z. Chen, *et al.*, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [23] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "Wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [24] T. Schultz, N. T. Vu, and T. Schlippe, "Globalphone: A multilingual text & speech database in 20 languages," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, IEEE, 2013, pp. 8126–8130.
- [25] R. Ardila, M. Branson, K. Davis, *et al.*, "Common Voice: A Massively-Multilingual Speech Corpus," in *Proceedings of the Twelfth Language Resources and Evaluation Conference*, N. Calzolari, F. Béchet, P. Blache, *et al.*, Eds., Marseille, France: European Language Resources Association, May 2020, pp. 4218–4222.
- [26] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "Sdr–half-baked or well done?" In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2019, pp. 626–630.
- [27] J. Cosentino, M. Pariente, S. Cornell, A. Deleforge, and E. Vincent, *Librimix: An open-source dataset for generalizable speech separation*, 2020. arXiv: 2005.11262 [eess.AS].
- [28] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *2016 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, IEEE, 2016, pp. 31–35.
- [29] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An asr corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5206–5210.
- [30] D. B. Paul and J. Baker, "The design for the wall street journal-based csr corpus," in *Speech and Natural Language: Proceedings of a Workshop Held at Harriman, New York, February 23-26, 1992*, 1992.
- [31] M. Ravanelli, T. Parcollet, P. Plantinga, *et al.*, *SpeechBrain: A general-purpose speech toolkit*, arXiv:2106.04624, 2021. arXiv: 2106.04624 [eess.AS].
- [32] M. Zanon Boito, V. Iyer, N. Lagos, L. Besacier, and I. Calapodescu, "mHuBERT-147: A Compact Multilingual HuBERT Model," in *Interspeech 2024*, ISCA, Sep. 1, 2024, pp. 3939–3943.
- [33] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, "Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs," in *2001 IEEE international conference on acoustics, speech, and signal processing. Proceedings (Cat. No. 01CH37221)*, IEEE, vol. 2, 2001, pp. 749–752.
- [34] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, "A short-time objective intelligibility measure for time-frequency weighted noisy speech," in *2010 IEEE international conference on acoustics, speech and signal processing*, IEEE, 2010, pp. 4214–4217.