

# TYPING TO LISTEN AT THE COCKTAIL PARTY: TEXT-GUIDED TARGET SPEAKER EXTRACTION

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Humans possess an extraordinary ability to selectively focus on the sound source of interest amidst complex acoustic environments, commonly referred to as cocktail party scenarios. In an attempt to replicate this remarkable auditory attention capability in machines, target speaker extraction (TSE) models have been developed. These models leverage the pre-registered cues of the target speaker to extract the sound source of interest. However, the effectiveness of these models is hindered in real-world scenarios due to the unreliable or even absence of pre-registered cues. To address this limitation, this study investigates the integration of natural language description to enhance the feasibility, controllability, and performance of existing TSE models. Specifically, we propose a model named LLM-TSE, wherein a large language model (LLM) extracts useful semantic cues from the user’s typed text input. These cues can serve as independent extraction cues, task selectors to control the TSE process or complement the pre-registered cues. Our experimental results demonstrate competitive performance when only text-based cues are presented, the effectiveness of using input text as a task selector, and a new state-of-the-art when combining text-based cues with pre-registered cues. To our knowledge, this is the first study to successfully incorporate LLMs to guide target speaker extraction, which can be a cornerstone for cocktail party problem research. Demos are provided at <https://github.com/LLM-TSE/llm-tse.github.io><sup>1</sup>

## 1 INTRODUCTION

The “Cocktail Party Problem” (E. Colin, 1953) – a term coined to describe a scenario where multiple sound sources are engaged in simultaneous conversation, yet a listener can selectively concentrate on a single sound source. This scenario represents a complex challenge in auditory perception (Haykin & Chen, 2005; Mesgarani & Chang, 2012; Bizley & Cohen, 2013) and serves as a remarkable demonstration of the intricate sound processing that occurs within the human auditory system. The human auditory system manages this complexity with remarkable efficacy, seemingly with ease. However, machines, such as hearing-aid devices (Shinn-Cunningham & Best, 2008), teleconferencing systems (Chen et al., 2020; Raj et al., 2021; Yoshioka et al., 2018), and hands-free human-machine interfaces (e.g., TVs, smartphones) (Gannot et al., 2017), encounter significant challenges in the context where multiple speakers talk at the same time.

Studies on computational auditory scene analysis (CASA) (Lyon, 1983; Meddis & Hewitt, 1991; Seltzer et al., 2003; Wang & Brown, 2006), non-negative matrix factorization (NMF) (Cichocki et al., 2006; Virtanen, 2007; Parry & Essa, 2007), and factorial Hidden Markov Models and Gaussian Mixture Models (HMM-GMM) (Virtanen, 2006; Stark et al., 2011) provide invaluable insights into solving the cocktail party problem. However, these methods are often limited by the representation power of their models, resulting in poor performance in complex acoustic environments. The advent of deep learning has paved the way for the application of deep neural networks (DNNs) in addressing this challenging problem. These existing DNN-based techniques can be broadly classified into two main categories: blind source separation (BSS) (Pal et al., 2013; Hershey et al., 2016; Yu et al., 2017; Luo & Mesgarani, 2019) and target speaker extraction (TSE) (Luo et al., 2018; Žmolíková et al., 2019; Xu et al., 2020; Ge et al., 2020; Pan et al., 2022; Zmolikova et al., 2023).

<sup>1</sup>Source code and datasets will be publicly available after review.




Figure 1: Comparison between the conventional TSE system and our proposed Text-Guided TSE system. The conventional systems rely on the pre-registered voiceprint of the target speaker as an extraction cue, while our system offers the flexibility to incorporate text-based cues to facilitate the target speaker extraction.

BSS techniques usually adopts DNNs to estimate an auditory mask for each speaker. The mask is then leveraged to separate each speaker’s voice into an individual stream from the mixture speech captured by a microphone. A difficulty in this process is the problem of global permutation ambiguity (Hershey et al., 2016), which occurs when attempting to accurately assign the output of a multi-source separation system to the correct source. To address this ambiguity problem, deep clustering (DC) techniques (Hershey et al., 2016; Isik et al., 2016; Wang et al., 2018) were proposed to group the spectro-temporal features belonging to the same speaker together through a clustering scheme. Permutation invariant training (PIT) (Yu et al., 2017; Kolbæk et al., 2017) was invented by finding the minimal loss over all the permutations between the extracted streams and the reference speeches. Typically, these methods require prior knowledge or estimation of the number of speakers in the mixture. However, in real-world scenarios, the number of speakers is hard to predict in advance.

Target speaker extraction provides an alternative solution to address the challenges of the unknown number of speakers and global permutation ambiguity. This approach involves providing a cue that is related to the desired speaker, such as a pre-recorded speech describing the voice characteristics (Xu et al., 2020), a spatial cue indicating the speaker’s direction (Ge et al., 2022), or synchronous lip movement (Pan et al., 2022). By using these specified cues, only the target speaker’s voice is extracted, thereby avoiding the issue of the unknown number of speakers and global permutation ambiguity. However, these pre-registered cues may vary substantially or even be absent in real environments, limiting the effectiveness of these systems.

To overcome the aforementioned limitation, as shown in Figure 1, we propose a novel text-guided TSE model, LLM-TSE, incorporating text descriptions as additional cues to enhance the feasibility, controllability, and performance of existing TSE models. Specifically, we leverage the power of large language models (LLMs) to extract meaningful semantic cues from the user’s typed text input. These text descriptions encompass various aspects of human auditory perception, including speaker characteristics, language, conversation contents, room characteristics, etc. These cues can serve as independent extraction cues, task selectors to control the TSE process or complement the pre-registered cues. By incorporating text descriptions as additional cues, we demonstrate that the performance of TSE models is significantly enhanced in various scenarios. The contributions of this work can be summarised as follows:

- To the best of our knowledge, this is the first study to utilize natural language description as extraction cues for target speaker extraction. We show these semantic cues possess high discriminative power and, therefore, can significantly enhance the feasibility of existing TSE methods.
- Our system implements a control mechanism through the natural language description to facilitate the speaker extraction process. This approach enables us to selectively retain or remove the source of interest based on the semantic concepts expressed in the text. By using text as a control mechanism, our system becomes a unified and flexible approach that eliminates the need for training multiple systems.
- Our system represents a significant advancement in TSE by integrating context-dependent information from typed descriptions with pre-registered cues. Unlike traditional cues, typically pre-recorded and isolated from the current acoustic environment, Our system captures complement cues from human perception. By incorporating additional cues that align with human

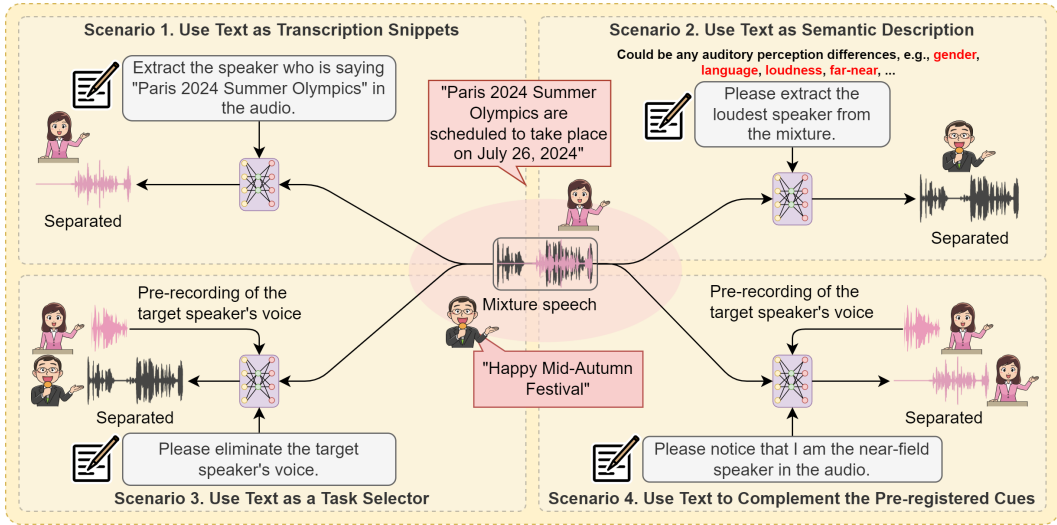


Figure 2: New application scenarios enabled by the proposed LLM-TSE model.

perception, our system achieves a more accurate and comprehensive representation of speech mixtures, thereby improving the effectiveness of TSE in practical scenarios.

## 2 TEXT-GUIDED TARGET SPEAKER EXTRACTION

The proposed LLM-TSE model opens up a plethora of novel application scenarios, surpassing the capabilities of traditional TSE techniques. As depicted in Figure 2, these application scenarios can be divided into the following four categories:

**Use text as transcription snippets:** Humans utilize discernible cues in relatively clean speech segments to enhance the perception of highly corrupted speech segments. Analogously, the LLM-TSE model can leverage distinguishable acoustic cues, in the form of transcription snippets, to facilitate speaker extraction, surpassing the capabilities of current TSE models.

**Use text as semantic description:** Apart from the above content-based cues, humans employ many other perceptual cues based on the distinguishing characteristics between competing speakers, such as gender, language, loudness level, and reverberation in the audio signal. The LLM-TSE model enables users to incorporate such perceptual cues as text-based semantic descriptions to exert control over the process of target speaker extraction. Notably, these perceptual cues can be considered as independent pre-registered cues.

**Use text as a task selector:** During a conversation involving multiple speakers, humans often switch their focus from one speaker to another. In addition, the speaker of interest at one moment may become a distraction at a later moment. In contrast to existing TSE systems that can only concentrate on a pre-registered speaker, the proposed LLM-TSE model empowers users with the flexibility to decide whether to retain or exclude the pre-registered speaker from the audio mixture, expanding the capabilities beyond what is currently achievable.

**Use text to complement the pre-registered cues.** In conventional TSE systems, the voice of the target speaker is typically pre-recorded in an acoustic environment that may differ substantially from the actual deployment environments. This discrepancy significantly affects the robustness of conventional TSE systems. In contrast, the proposed LLM-TTS model has the ability to compensate for these differences by providing complementary cues in addition to the pre-registered ones, such as the speaker’s location, language, loudness level, etc. Consequently, it generates a more comprehensive and accurate representation of the target speaker that can significantly enhance the system’s robustness.

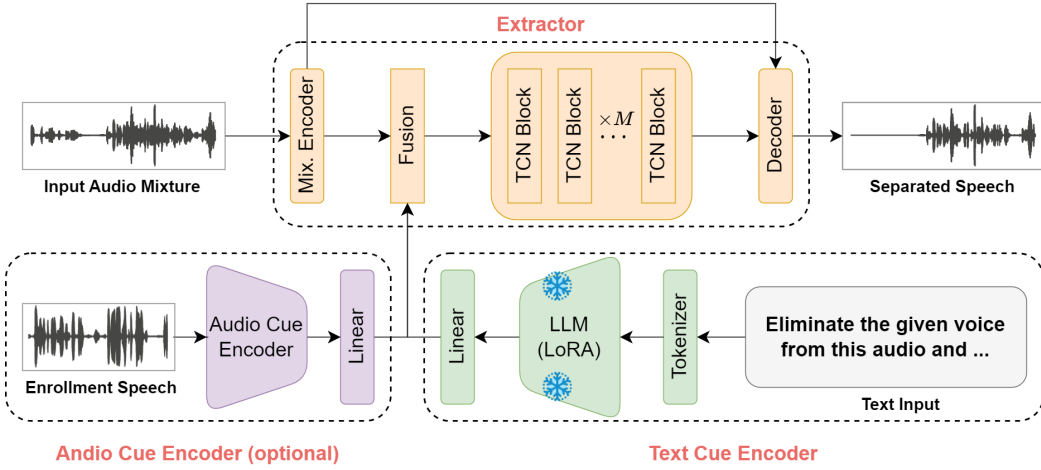


Figure 3: Overview of the proposed LLM-TSE model architecture.

### 3 LLM-TSE MODEL

As illustrated in Figure 3, the proposed LLM-TSE model follows a processing pipeline of Encoding-Fusion-Extraction-Decoding. In the encoding phase, three distinct encoders are employed to convert the pre-registered speech, text prompts, and input audio mixture into corresponding embeddings. Leveraging the fused embeddings representing the enrolled speech and text cues, the extractor then selectively extracts the desired sound source from the input audio mixture. Finally, the frequency-domain feature representation obtained from the extractor is transformed back into the time-domain and output as the extracted speech.

**Mixture Encoder and Decoder:** The mixture encoder transforms the input audio mixture from the time domain into the feature representation, which can be more effectively handled by the extractor. This transformation is realized by convolving each audio frame of length  $L$  with a set of  $N$  1-D convolution filters  $\{u_n(t)\}_{n=\{0\dots N-1\}}$ , which can be expressed as follows:

$$\mathbf{X}(k, n) = \sum_{t=0}^{L-1} x(t + kH)u_n(t), \quad n \in \{0, \dots, N-1\}, \quad (1)$$

where  $x(t)$  is the input signal,  $k \in \{0, \dots, K-1\}$  is the frame index,  $H$  is the hop size, and  $\mathbf{X}(k, n)$  is the result of the convolution operation. Similarly, the decoder maps the extracted feature, denoted as  $\mathbf{Y}(k, n)$ , back to the time domain via a transposed 1-D convolution operation with  $N$  synthesis filters  $\{v_n(t)\}_{n=\{0\dots N-1\}}$ , and each has a length of  $L$ :

$$\hat{y}(t) = \sum_{k=0}^{K-1} \sum_{n=0}^{N-1} \mathbf{Y}(k, n)v_n(t - kH), \quad (2)$$

where  $\hat{y}(t)$  is the extracted audio signal in time domain.

**Text Cue Encoder:** We utilize the LLaMA-2 7B Chat LLM, a dialogue-fine-tuned version of the LLaMA-2 (Touvron et al., 2023), to obtain discriminative semantic embeddings from the user’s text input. LLaMA-2 is pre-trained on a combination of natural language and programming language corpora in a self-supervised manner. LLaMA-2 7B Chat LLM is further fine-tuned from LLaMA-2 via instruction-tuning, which significantly enhances its performance on various reasoning and generation tasks. During our model training, instead of performing full fine-tuning on the adopted LLM text encoder, we adopt the parameter-efficient Low-Rank Adaptation (LoRA) technique (Hu et al., 2021). LoRA introduces a small set of parameters into the frozen LLaMA-2 7B Chat LLM, which are referred to as LoRA adapters. Specifically, one LoRA adapter is attached to each LLM layer, modifying its frozen parameter by adding a low-rank learnable matrix of the same size. In the proposed LLM-TSE model, we apply the LoRA adapters to only modify keys and queries in each self-attention layer. Ultimately, we only add 12% more trainable parameters. This approach

not only helps to prevent the overfitting problem that is often encountered with a small fine-tuning dataset but also improves the training efficiency.

**Audio Cue Encoder:** The primary role of the audio cue encoder is to encode the optional pre-registered speech into a discriminative speaker embedding. The first step in this process involves transforming the time domain input signal, using the above-mentioned learnable 1-D convolutional filters, into the frequency domain. Following this transformation, we utilize a series of Temporal Convolutional Network (TCN) blocks (Pandey & Wang, 2019; Luo & Mesgarani, 2019) to extract speaker-related feature representation. These TCN blocks are designed to capture the temporal dependencies in the speech signal, which are crucial for distinguishing different speakers. Finally, we take the average along the temporal dimension to generate a speaker embedding vector, which effectively captures the unique vocal attributes of the pre-registered speech that can differentiate one speaker from others.

**Fusion Layer:** Here, we follow a simple concatenation approach to fuse the audio and text cues, which has shown to be effective in many other TSE systems (Žmolíková et al., 2019; Ge et al., 2020). Specifically, we transform the text cue and audio cue embeddings into the same dimensionality through two linear projection layers, and then directly concatenate them to form a multi-modal representation.

**Extractor:** The last part of our model is the target extractor, which serves to estimate the target signal. We adopt the widely used time-frequency masking-based extractor (Luo & Mesgarani, 2019; Isik et al., 2016), whose operations can be summarized as follows:

$$\begin{aligned} \mathbf{M} &= \text{MaskNet}(\mathbf{Z}; \theta^{\text{Mask}}), \\ \hat{\mathbf{X}} &= \mathbf{M} \otimes \mathbf{X}, \end{aligned} \tag{3}$$

where  $\mathbf{Z}$  is the fused embedding generated from the fusion layer,  $\text{MaskNet}(\cdot)$  is a TCN-based NN that estimates the time-frequency mask  $\mathbf{M} \in \mathbb{R}^{D \times N}$  for the target speaker, where  $D$  is the feature dimension of each time step.  $\theta^{\text{Mask}}$  is the network parameter, and  $\otimes$  denotes the element-wise Hadamard product.  $\hat{\mathbf{X}}$  is the estimated target speech signal in the frequency domain.

**Loss function:** The parameters of the proposed LLM-TSE model are optimized by minimizing the following Scale-Invariant Signal-to-Distortion Ratio (SI-SDR) (Roux et al., 2019) loss function:

$$\mathcal{L}^{\text{SI-SDR}} = -10 \log_{10} \left( \frac{\|\frac{\hat{\mathbf{y}}^T \mathbf{y}}{\|\mathbf{y}\|^2} \mathbf{y}\|^2}{\|\frac{\hat{\mathbf{y}}^T \mathbf{y}}{\|\mathbf{y}\|^2} \mathbf{y} - \hat{\mathbf{y}}\|^2} \right). \tag{4}$$

The SI-SDR loss is computed directly in the time domain, which forces the model to learn how to precisely estimate the magnitude and the phase of the target speech signals.

## 4 EXPERIMENTAL EVALUATION

In this paper, our primary objective is to integrate text-based cues to enhance the target speaker extraction systems. In the following sections, we initially delve into the method of simulating the overlapped mixture of speech data. Subsequently, we will explore the generation of text questions.

### 4.1 OVERLAPPED SPEECH SIMULATION

Our experiment uses two speech datasets: LibriSpeech (Panayotov et al., 2015) and Multilingual LibriSpeech (MLS) (Pratap et al., 2020). LibriSpeech, a 1000-hour corpus of English audiobook speech, is known for its diverse speaker identities. MLS, an extension of LibriSpeech, adds multiple languages, including French, German, Spanish, etc. Due to it having too much data, we randomly selected 400 speakers per language from MLS with up to 20 utterances each. We adhered to LibriSpeech’s standard training, validation, and test set division. For MLS, we randomly assigned 5% of speakers from each language to validation and test sets, respectively, with the rest for training.

Our experiments cover a variety of attributes, including transcription snippets, gender, language, loudness, and far-near. For transcription snippets extraction, we only use the LibriSpeech dataset

and the corresponding pre-extracted forced alignment (Chodroff, 2023) data<sup>2</sup> to identify the word timestamps from LibriSpeech. The remainder of the data for simulation is randomly selected from the LibriSpeech and MLS datasets. For generating the mixture speech, we adopt online simulation, generating the data needed for each iteration beforehand. The number of speakers in the mixture of speech is limited to two, stipulating that the two speakers have different attributes for gender, language, loudness, or far-near. When generating a mixture of speech for the loudness task, our signal-to-noise ratio is randomly selected from -3 dB to -2 dB and 2 dB to 3 dB. The other tasks span from -3 dB to 3 dB. In the case of the distance task, we include both near (target speaker) - far (interference speaker) and far (interference speaker) - near (target speaker) scenarios. For the other tasks, near and far combinations are randomized. Room dimensions are randomly selected from lengths of 9 to 11 m, widths of 9 to 11 m, and heights of 2.6 to 3.5 m. The reverberation time ranges from 0.3 to 0.6 seconds. We use Pyroomacoustics<sup>3</sup> to generate Room Impulse Responses (RIRs), and the microphone’s position is defaulted to the center of the room. The sound source distance from the microphone varies between 0.3 to 0.5 m and 1.5 to 2.5 m for near or far fields, respectively. The angle ranges from 0 to 180 degree, and the sound source’s height varies between 1.6 to 1.9 m.

The mixture and pre-registered speeches are set to a duration of 6 seconds, with a randomly determined overlap ratio between 40% and 70%. The pre-registered speech is randomly selected from the remaining target speaker’s speech. If the training objective is to remove the target speaker, the other speaker’s speech from the mixture serves as the training target. We assume that each generated mixture speech sample should exhibit a distinguishable attribute throughout the training. All experimental data is sampled at 16,000 Hz to ensure high-quality audio.

## 4.2 TEXT GENERATION

We include three types of texts to explore using LLMs to enrich target speaker extraction systems. We first create ten foundational question templates for each type of task. These templates will then be rephrased and expanded using ChatGPT-4-32K<sup>4</sup> to produce 100 diverse text prompts. We adopt a non-overlapped 80/10/10% partitioning for training, validation, and testing sets. The text prompts used in the testing set are unseen during the training.

**Text as an independent extraction cue:** In this type, the text is used as an independent extraction cue. The texts of this task are like: “Extracting a voice with  $\langle$  specific characteristic  $\rangle$  from a mixture of speech”, e.g., scenarios 1&2 in Figure 2. The text description outlines the features of the voice to be extracted, including the transcription snippets of the mixture of speech, the speaker’s language, gender, loudness, and far-near. For the transcription snippet task, we used 100% of the target speech text length as cues for training, testing with 50%, 80%, and 100% of the target speech text length to evaluate generalizability. This setup is highly functional, i.e., by informing the system about the audible part of the speech, the system can utilize both semantic and acoustic information to track and extract the desired speaker. It’s crucial to note that the attributes utilized in this study are not exhaustive. In real-world situations, humans employ a variety of other cues, e.g., emotion, pitch, etc., to extract the sound source of interest (Haykin & Chen, 2005; Shinn-Cunningham & Best, 2008). However, exploring these additional cues extends beyond the scope of this current study and is reserved for future research.

**Text as a task selector:** We propose one task type where texts can influence the system’s output: target speaker extraction or removal. The text serves as a directive for the system to either extract a given speaker’s voice or remove it from the mixture of audio. The generated texts are like “please remove the given voice from this audio.”

**Text as a complement to human perception in the audio-based extraction system:** We integrate the human understanding and interpretation of the mixture of speech into the extraction process, which can significantly enhance the system’s performance. Here, we cover all semantic types mentioned above, i.e., transcription snippets, gender, language, loudness, and far-near. The generated questions are like “Extracting a speaker based on the given pre-registered speech, where the speaker possesses a  $\langle$ specific characteristic $\rangle$  within the mixture speech.”

<sup>2</sup><https://github.com/CoirentinJ/librispeech-alignments>

<sup>3</sup><https://github.com/LCAV/pyroomacoustics>

<sup>4</sup><https://platform.openai.com/docs/models>

Table 1: Evaluation of SI-SDR (dB  $\uparrow$ ) metric across different methods. For the transcription snippet task, we use 100% of the target speech text as cues during training and test the model with a different amount of text transcriptions, including 50%, 80%, and 100%.

Entry	Inputs		Transcription Snippet			Gender	Language	Far-near	Loudness
	Audio	Text	50%	80%	100%				
Unproc.	-	-	-0.02	-0.02	-0.02	-0.02	-0.03	-0.01	-0.10
TD-SpeakerBeam	✓	✗	7.21			10.15	8.38	9.38	7.57
	✓	✗	7.30			10.17	8.87	9.77	7.75
LLM-TSE (LoRA Adapters, LLaMA-2 7B Chat)	✗	One-Hot	No Support			10.54	8.88	10.25	8.96
	✗	✓	2.70	3.97	7.48	10.40	9.38	10.57	8.89
	✓	One-Hot	No Support			10.62	10.18	10.32	8.99
	✓	✓	7.96	9.81	10.05	10.87	9.72	10.66	9.41
No LoRA Adapters (only Linear Projection)	✗	✓	1.66	3.38	5.38	8.76	7.38	8.45	5.46
	✓	✓	4.85	7.60	7.98	9.02	7.97	8.67	7.11
Use Vicuna-7b-v1.3 (Zheng et al. (2023))	✗	✓	2.23	3.31	8.79	9.44	8.29	9.27	5.75
	✓	✓	7.41	9.05	9.35	10.15	9.01	9.94	6.47

### 4.3 RESULTS

**Efficacy of Using Input Text as Independent Cues:** Table 1 demonstrates a notable performance enhancement when text alone is employed as an extraction cue, compared to unprocessed mixture speech. The proposed LLM-TSE model is built on TD-SpeakerBeam (Delcroix et al., 2020), a state-of-the-art (SOTA) open-source target speaker extraction model. Compared to TD-SpeakerBeam, the only modification in the LLM-TSE model is the additional text-prompt encoder. This enhancement is further corroborated by Figure 4. These findings suggest that the LLM-TSE model effectively interprets the provided text descriptions, which fundamentally serve as human interpretations of auditory object differences within a speech mixture. This innovative strategy represents a significant leap in harnessing natural language processing techniques for complex auditory tasks, thereby enhancing the scope of potential applications for speaker extraction methodologies.

**Efficacy of Using Input Text as Task Selector:** In this task, our objective is to control the training targets of the separation system using natural language. The corresponding textual queries could resemble “Is there a way to remove the given voice from this mixture audio?” In Figure 4, we illustrate the capacity of our system to determine whether to extract or suppress the sound source corresponding to the provided pre-registered speech when using text descriptions. Notably, the samples displayed in the third row exemplify this capability, as they successfully suppress the target sound source associated with the pre-registered speech. Our explorations in this area are somewhat limited at this stage. More broadly, these controls could be configured with greater flexibility. For instance, they could manipulate the degree of reverberation in the extracted speech (since individual preferences for reverberation vary) or dictate the impact range of the separation system (to avoid unnecessary non-linear-processing distortion). We intend to delve deeper into these aspects in our future work.

**Efficacy of Using Input Text to Complement the Pre-registered Cues:** Pre-registered speech primarily only encodes the speaker’s vocal characteristics regardless of any time or acoustic environmental context. We aim to introduce this contextual information into the target speaker extraction system utilizing text descriptions. For this purpose, a typical text description is like: “Separate the target speaker’s audio based on the provided pre-registered speech as a reference, bearing in mind that I am the speaker who employs a louder tone in the mixed speech.” The relevant experimental outcomes are presented in the middle section of Table 1. Upon integrating descriptions delineating auditory object differences, we noted a significant improvement in the system’s performance. This enhancement was particularly prominent in the “loudness” task, where the dataset contained a pronounced loudness disparity between the two sound sources. The challenge posed by identifying the target speaker using only the pre-registered speech was substantially mitigated upon implementing our approach, producing the most substantial performance increase within this task.

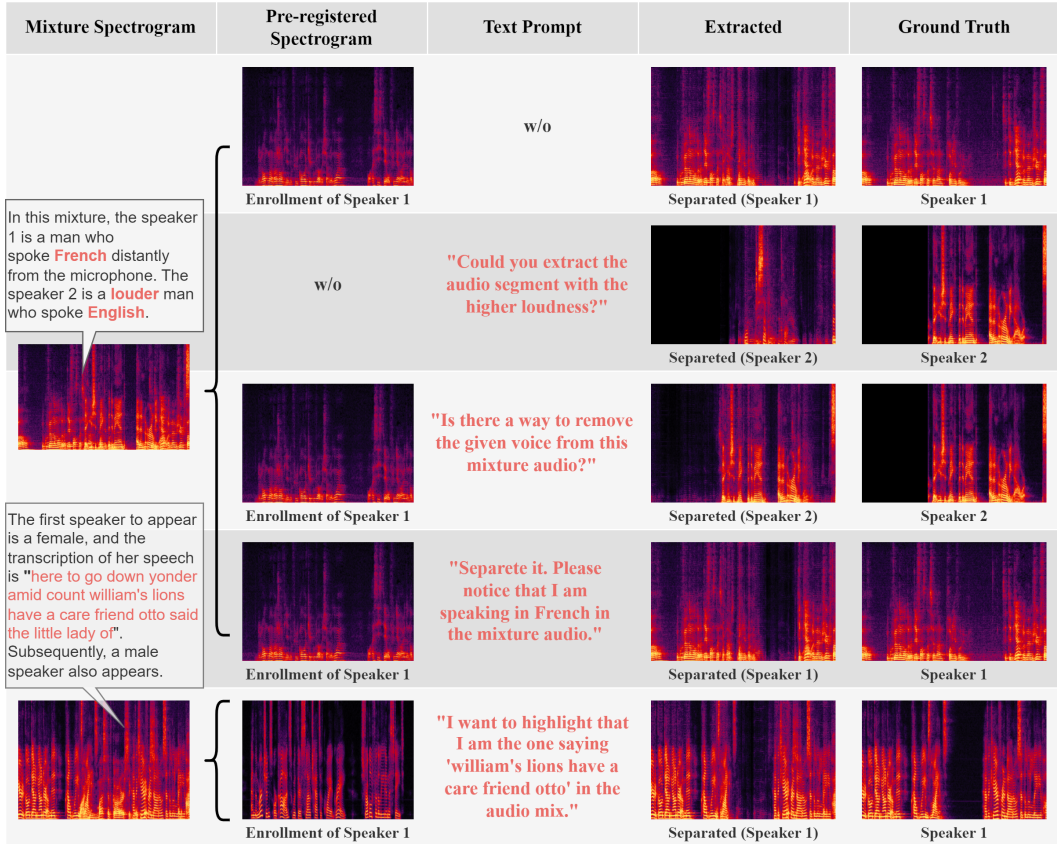


Figure 4: Samples generated from the proposed LLM-TSE model. The text box contains information about the input audio mixture. The term “w/o” indicates the absence of a certain input.

**Ablation Studies on Text Encoder Selection:** Here, we present the results of a sequence of ablation experiments executed on the text encoder component. The outcomes are summarized at the bottom of Table 1. At the outset, we assessed the functionality of the text cue encoder in the absence of the LoRA adaptors, where only the projection layer of the LLM model was permitted to train, effectively freezing all other parameters of the LLM. This configuration aimed to determine if the LLM’s generic understanding of diverse text corpora could offer sufficient discriminative information. However, our findings suggest that relying solely on embeddings, which are derived from the LLM’s interpretation of various text descriptions, is insufficient to accomplish the task whether an audio encoder was integrated into the system or not. In subsequent experiments, we employed the Vicuna 7B model (Zheng et al., 2023) as our text encoder. This model, which was fine-tuned on data from “shareGPT.com” and based on the LLaMA-v1 model, exhibited marginally inferior performance in natural language benchmark tasks compared to the LLAMA-2 7B Chat. Further, the Vicuna model underperformed in our target speaker separation task compared to the LLAMA-2 7B Chat. This observation supports the premise that employing a more powerful LLM as a text cue encoder can significantly enhance the discriminative capabilities of the overall system.

## 5 RELATED WORKS

**Audio-Language Multimodal Model:** Audio-language multimodal currently represents a significant research area with many application scenarios (Huang et al., 2023c; Zhang et al., 2023; Gong et al., 2023). The primary focus has revolved around audio events, with most tasks and datasets originating from automatic audio caption (Drossos et al., 2017; Wu et al., 2019; Mei et al., 2022), which aims to assign meaningful textual descriptions to audio content. Leveraging these datasets, related studies have been conducted on synthesizing audio based on text descriptions, which find applications in diverse scenarios such as film production, game design, and more. Among these, the Contrastive Language-Audio Pretraining (CLAP) (Elizalde et al., 2022) model is a large-scale



pre-training model that employs a contrastive learning approach similar to the Contrastive Language-Image Pretraining (CLIP) (Radford et al., 2021) model for aligning text and audio modalities. This model has pushed the boundaries in tasks that involve synthesizing audio based on text descriptions (Huang et al., 2023b; Kreuk et al., 2023; Liu et al., 2023a;b). Furthermore, the works conducted by Wang et al. (2023); Zhang et al. (2023); Le et al. (2023) expands the input modality to encompass audio and text instead of text only for audio generation. However, it’s important to note that the underlying logic is based on generative models that take audio and specific control inputs to handle various speech transformation tasks. These works are more like controlled speech/audio/music synthesis, not requiring the length of input and output to be strictly aligned. This is entirely different from the field of our study.

**Audio-Language-Vison Multimodal Target Source Separation:** Among all these audio-language multimodal models, those most relevant to our research involve separating or detecting audio events based on text description (Kilgour et al., 2022; Liu et al., 2022; 2023c; Li et al., 2023). These studies employ models like BERT (Devlin et al., 2019) (mini), CLAP, or other pre-trained models to comprehend descriptions of sound events, subsequently separating the sound sources consistent with the target description. However, they are not specifically designed for speech signals. In contrast to audio event classes, speech signals are considerably similar when observed from spectrograms, lacking clear acoustic spectral patterns to follow. Instead, they rely more on perceptual differences in auditory objects and semantic information. In addition to sound events, these models also focus on separating musical instruments (Chen et al., 2023; Huang et al., 2023a; Chen et al., 2023). It’s important to note that while these previous works have made significant strides in the field, the specific challenges and nuances of speech signal separation present a unique problem space that our work aims to address. Labels, particularly those implemented via one-hot vectors, can be seen as a distinctive type of human language. In the realm of label-based audio/music/speech extraction systems (Manilow et al., 2020; Delcroix et al., 2021; Tzinis et al., 2022; Delcroix et al., 2023; Li et al., 2023; Ochiai et al., 2020), the works of Manilow et al. (2020) and Tzinis et al. (2022) are most closely aligned with ours. These systems, like ours, endeavor to integrate human subjective intentions into the separation process through attribute labels. Yet, they solely rely on one-hot vectors, resulting in a lack of flexibility within human-computer dialogue systems. In addition, they cannot understand the vast array of human language inputs and struggle significantly when dealing with open-ended queries. By contrast, we employ LLMs to understand cues that extend beyond human descriptions of auditory object differences, which offers increased flexibility in cue extraction. Furthermore, we’ve investigated control capabilities and made a connection between the perceptual differences of auditory objects in mixture and voiceprint systems. Another method utilizes semantic cues, such as images Ohishi et al. (2022), to extract speakers’ speech discussing a particular concept. However, the necessity for corresponding images constrains its potential application domains.

## 6 CONCLUSION AND FUTURE WORKS

In this study, we proposed a novel paradigm for target speaker extraction, namely LLM-TSE, a significant departure from previous methodologies. Our approach uniquely introduces text to provide useful speaker extraction cues, which is an innovation that has demonstrated notable success and improvement in our experimental results. Our investigations have illuminated the potential of natural language to provide a rich source of discriminative features. These features can be leveraged independently as extraction cues, showcasing the versatility and effectiveness of natural language in this context. Furthermore, natural language is useful for performing task selection, which represents a promising approach to achieving auditory attention switching. Moreover, our paradigm augments the performance of audio-only systems by integrating contextual information from the present acoustic environment, which is often overlooked in traditional methods. This addition provides a more comprehensive and accurate representation of the target speaker’s context, further enhancing the extraction process. In summary, our proposed paradigm signifies an important advancement for target speaker extraction systems, extending accessibility and improving performance. Not only does it provide a fresh perspective on the extraction process, but it also lays the groundwork for potential future studies on the cocktail party problem. Moving forward, we plan to persist in this direction, enhancing machines’ ability to understand the foundations of human perception of multiple auditory objects within complex acoustic environments using natural language cues. Specifically, we aim to incorporate a range of mutually exclusive or non-exclusive auditory attributes, label flexible and open-ended text descriptions, and develop the capability for multi-round separation.

## REFERENCES

- Jennifer K. Bizley and Yale E. Cohen. The what, where and how of auditory-object perception. *Nature Reviews Neuroscience*, 14(10):693–707, October 2013. ISSN 1471-0048. 301 citations (Crossref) [2023-06-21] 346 citations (Semantic Scholar/DOI) [2023-06-21] Number: 10 Publisher: Nature Publishing Group.
- Ke Chen, Yusong Wu, Haohe Liu, Marianna Nezhurina, Taylor Berg-Kirkpatrick, and Shlomo Dubnov. MusicLDM: Enhancing Novelty in Text-to-Music Generation Using Beat-Synchronous Mixup Strategies, August 2023. arXiv:2308.01546 [cs, eess].
- Zhuo Chen, Takuya Yoshioka, Liang Lu, Tianyan Zhou, Zhong Meng, Yi Luo, Jian Wu, Xiong Xiao, and Jinyu Li. Continuous Speech Separation: Dataset and Analysis. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7284–7288, May 2020. ISSN: 2379-190X.
- Eleanor Chodroff. Montreal Forced Aligner. *Linguistics Methods Hub*, January 2023.
- Andrzej Cichocki, Rafal Zdunek, and Shun-ichi Amari. New algorithms for non-negative matrix factorization in applications to blind source separation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 5, pp. V–V, 2006.
- Marc Delcroix, Tsubasa Ochiai, Katerina Zmolikova, Keisuke Kinoshita, Naohiro Tawara, Tomohiro Nakatani, and Shoko Araki. Improving speaker discrimination of target speech extraction with time-domain SpeakerBeam, January 2020. 60 citations (Semantic Scholar/arXiv) [2023-02-14] arXiv:2001.08378 [cs, eess].
- Marc Delcroix, Jorge Bennisar Vázquez, Tsubasa Ochiai, Keisuke Kinoshita, and Shoko Araki. Few-Shot Learning of New Sound Classes for Target Sound Extraction. In *Interspeech 2021*, pp. 3500–3504. ISCA, August 2021.
- Marc Delcroix, Jorge Bennisar Vázquez, Tsubasa Ochiai, Keisuke Kinoshita, Yasunori Ohishi, and Shoko Araki. SoundBeam: Target Sound Extraction Conditioned on Sound-Class Labels and Enrollment Clues for Increased Performance and Continuous Learning. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:121–136, 2023. ISSN 2329-9304. Conference Name: IEEE/ACM Transactions on Audio, Speech, and Language Processing.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.
- Konstantinos Drossos, Sharath Adavanne, and Tuomas Virtanen. Automated audio captioning with recurrent neural networks. In *2017 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, pp. 374–378, October 2017. ISSN: 1947-1629.
- Cherry E. Colin. Some experiments on the recognition of speech, with one and with two ears. *The Journal of The Acoustical Society of America*, 25(5):975–979, 1953.
- Benjamin Elizalde, Soham Deshmukh, Mahmoud Al Ismail, and Huaming Wang. CLAP: Learning Audio Concepts From Natural Language Supervision, June 2022. arXiv:2206.04769 [cs, eess].
- Sharon Gannot, Emmanuel Vincent, Shmulik Markovich-Golan, and Alexey Ozerov. A Consolidated Perspective on Multimicrophone Speech Enhancement and Source Separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 25(4):692–730, April 2017. ISSN 2329-9304. Conference Name: IEEE/ACM Transactions on Audio, Speech, and Language Processing.
- Meng Ge, Chenglin Xu, Longbiao Wang, Eng Siong Chng, Jianwu Dang, and Haizhou Li. SpEx+: A Complete Time Domain Speaker Extraction Network. In *Interspeech 2020*, pp. 1406–1410. ISCA, October 2020.

- Meng Ge, Chenglin Xu, Longbiao Wang, Eng Siong Chng, Jianwu Dang, and Haizhou Li. L-spex: Localized target speaker extraction. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7287–7291. IEEE, 2022.
- Yuan Gong, Hongyin Luo, Alexander H. Liu, Leonid Karlinsky, and James Glass. Listen, Think, and Understand, May 2023. arXiv:2305.10790 [cs, eess] version: 1.
- Simon Haykin and Zhe Chen. The Cocktail Party Problem. *Neural Computation*, 17(9):1875–1902, September 2005. ISSN 0899-7667.
- John R. Hershey, Zhuo Chen, Jonathan Le Roux, and Shinji Watanabe. Deep clustering: Discriminative embeddings for segmentation and separation. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 31–35, Shanghai, March 2016. IEEE. ISBN 978-1-4799-9988-0.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-Rank Adaptation of Large Language Models, October 2021. arXiv:2106.09685 [cs].
- Qingqing Huang, Daniel S. Park, Tao Wang, Timo I. Denk, Andy Ly, Nanxin Chen, Zhengdong Zhang, Zhishuai Zhang, Jiahui Yu, Christian Frank, Jesse Engel, Quoc V. Le, William Chan, Zhifeng Chen, and Wei Han. Noise2Music: Text-conditioned Music Generation with Diffusion Models, March 2023a. arXiv:2302.03917 [cs, eess].
- Rongjie Huang, Jiawei Huang, Dongchao Yang, Yi Ren, Luping Liu, Mingze Li, Zhenhui Ye, Jinglin Liu, Xiang Yin, and Zhou Zhao. Make-An-Audio: Text-To-Audio Generation with Prompt-Enhanced Diffusion Models, January 2023b. arXiv:2301.12661 [cs, eess].
- Rongjie Huang, Mingze Li, Dongchao Yang, Jiatong Shi, Xuankai Chang, Zhenhui Ye, Yuning Wu, Zhiqing Hong, Jiawei Huang, Jinglin Liu, Yi Ren, Zhou Zhao, and Shinji Watanabe. AudioGPT: Understanding and Generating Speech, Music, Sound, and Talking Head, April 2023c. arXiv:2304.12995 [cs, eess].
- Yusuf Isik, Jonathan Le Roux, Zhuo Chen, Shinji Watanabe, and John R. Hershey. Single-Channel Multi-Speaker Separation Using Deep Clustering. In *Interspeech 2016*, pp. 545–549. ISCA, September 2016.
- Kevin Kilgour, Beat Gfeller, Qingqing Huang, Aren Jansen, Scott Wisdom, and Marco Tagliasacchi. Text-Driven Separation of Arbitrary Sounds, April 2022. arXiv:2204.05738 [cs, eess].
- Morten Kolbæk, Dong Yu, Zheng-Hua Tan, and Jesper Jensen. Multitalker Speech Separation With Utterance-Level Permutation Invariant Training of Deep Recurrent Neural Networks. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 25(10):1901–1913, October 2017. ISSN 2329-9304. Conference Name: IEEE/ACM Transactions on Audio, Speech, and Language Processing.
- Felix Kreuk, Gabriel Synnaeve, Adam Polyak, Uriel Singer, Alexandre Défossez, Jade Copet, Devi Parikh, Yaniv Taigman, and Yossi Adi. AudioGen: Textually Guided Audio Generation, March 2023. arXiv:2209.15352 [cs, eess].
- Matthew Le, Apoorv Vyas, Bowen Shi, Brian Karrer, Leda Sari, Rashel Moritz, Mary Williamson, Vimal Manohar, Yossi Adi, Jay Mahadeokar, and Wei-Ning Hsu. Voicebox: Text-Guided Multilingual Universal Speech Generation at Scale, June 2023. arXiv:2306.15687 [cs, eess].
- Chenda Li, Yao Qian, Zhuo Chen, Dongmei Wang, Takuya Yoshioka, Shujie Liu, Yanmin Qian, and Michael Zeng. Target Sound Extraction with Variable Cross-Modality Clues. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5, Rhodes Island, Greece, June 2023. IEEE. ISBN 978-1-72816-327-7.
- Haohe Liu, Zehua Chen, Yi Yuan, Xinhao Mei, Xubo Liu, Danilo Mandic, Wenwu Wang, and Mark D. Plumbley. AudioLDM: Text-to-Audio Generation with Latent Diffusion Models. In *Proceedings of the 40th International Conference on Machine Learning*, pp. 21450–21474. PMLR, July 2023a. ISSN: 2640-3498.

- Haohe Liu, Qiao Tian, Yi Yuan, Xubo Liu, Xinhao Mei, Qiuqiang Kong, Yuping Wang, Wenwu Wang, Yuxuan Wang, and Mark D. Plumbley. AudioLDM 2: Learning Holistic Audio Generation with Self-supervised Pretraining, September 2023b. arXiv:2308.05734 [cs, eess].
- 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. In *Interspeech 2022*, pp. 1801–1805. ISCA, September 2022.
- 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, August 2023c. arXiv:2308.05037 [cs, eess].
- 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*, 27(8):1256–1266, August 2019. ISSN 2329-9290, 2329-9304.
- Yi Luo, Zhuo Chen, and Nima Mesgarani. Speaker-Independent Speech Separation With Deep Attractor Network. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 26(4):787–796, April 2018. ISSN 2329-9290, 2329-9304.
- Richard Lyon. A computational model of binaural localization and separation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 8, pp. 1148–1151, 1983.
- Ethan Manilow, Gordon Wichern, and Jonathan Roux Le. Hierarchical Musical Instrument Separation. In *International Society for Music Information Retrieval Conference*, pp. 1–8, VIRTUAL CONFERENCE, October 2020.
- Ray Meddis and Michael J Hewitt. Virtual pitch and phase sensitivity of a computer model of the auditory periphery. i: Pitch identification. *The Journal of the Acoustical Society of America*, 89(6):2866–2882, 1991.
- Xinhao Mei, Xubo Liu, Mark D. Plumbley, and Wenwu Wang. Automated audio captioning: an overview of recent progress and new challenges. *EURASIP Journal on Audio, Speech, and Music Processing*, 2022(1):26, October 2022. ISSN 1687-4722.
- Nima Mesgarani and Edward F. Chang. Selective cortical representation of attended speaker in multi-talker speech perception. *Nature*, 485(7397):233–236, May 2012. ISSN 1476-4687. 790 citations (Semantic Scholar/DOI) [2023-09-26] Number: 7397 Publisher: Nature Publishing Group.
- Tsubasa Ochiai, Marc Delcroix, Yuma Koizumi, Hiroaki Ito, Keisuke Kinoshita, and Shoko Araki. Listen to What You Want: Neural Network-Based Universal Sound Selector. In *Interspeech 2020*, pp. 1441–1445. ISCA, October 2020.
- Yasunori Ohishi, Marc Delcroix, Tsubasa Ochiai, Shoko Araki, Daiki Takeuchi, Daisuke Niizumi, Akisato Kimura, Noboru Harada, and Kunio Kashino. ConceptBeam: Concept Driven Target Speech Extraction. In *Proceedings of the 30th ACM International Conference on Multimedia*, pp. 4252–4260, Lisboa Portugal, October 2022. ACM. ISBN 978-1-4503-9203-7.
- Madhab Pal, Rajib Roy, Joyanta Basu, and Milton S. Bepari. Blind source separation: A review and analysis. In *2013 International Conference Oriental COCODSA held jointly with 2013 Conference on Asian Spoken Language Research and Evaluation (O-COCODSA/CASLRE)*, pp. 1–5, Gurgaon, India, November 2013. IEEE. ISBN 978-1-4799-2378-6.
- Zexu Pan, Ruijie Tao, Chenglin Xu, and Haizhou Li. Selective listening by synchronizing speech with lips. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30:1650–1664, 2022.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: An ASR corpus based on public domain audio books. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5206–5210, South Brisbane, Queensland, Australia, April 2015. IEEE. ISBN 978-1-4673-6997-8.

- Ashutosh Pandey and DeLiang Wang. TCNN: Temporal Convolutional Neural Network for Real-time Speech Enhancement in the Time Domain. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6875–6879, Brighton, United Kingdom, May 2019. IEEE. ISBN 978-1-4799-8131-1.
- R Mitchell Parry and Irfan Essa. Incorporating phase information for source separation via spectrogram factorization. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 2, pp. II–661, 2007.
- Gerald R. Popelka, Brian C. J. Moore, Richard R. Fay, and Arthur N. Popper. *Hearing Aids*, volume 56 of *Springer Handbook of Auditory Research*. Springer International Publishing, Cham, 2016. ISBN 978-3-319-33034-1 978-3-319-33036-5.
- Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, and Ronan Collobert. MLS: A Large-Scale Multilingual Dataset for Speech Research. In *Interspeech 2020*, pp. 2757–2761, October 2020. arXiv:2012.03411 [cs, eess].
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 8748–8763. PMLR, July 2021. ISSN: 2640-3498.
- Desh Raj, Pavel Densov, Zhuo Chen, Hakan Erdogan, Zili Huang, Maokui He, Shinji Watanabe, Jun Du, Takuya Yoshioka, Yi Luo, Naoyuki Kanda, Jinyu Li, Scott Wisdom, and John R. Hershey. Integration of Speech Separation, Diarization, and Recognition for Multi-Speaker Meetings: System Description, Comparison, and Analysis. In *2021 IEEE Spoken Language Technology Workshop (SLT)*, pp. 897–904, January 2021.
- Jonathan Le Roux, Scott Wisdom, Hakan Erdogan, and John R. Hershey. SDR – Half-baked or Well Done? In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 626–630, May 2019. ISSN: 2379-190X, 1520-6149.
- Michael L Seltzer, Jasha Droppo, and Alex Acero. A harmonic-model-based front end for robust speech recognition. In *Eighth European Conference on Speech Communication and Technology*, 2003.
- Barbara G. Shinn-Cunningham and Virginia Best. Selective Attention in Normal and Impaired Hearing. *Trends in Amplification*, 12(4):283–299, 2008. ISSN 1084-7138. 356 citations (Semantic Scholar/DOI) [2023-09-26] 266 citations (Crossref) [2023-06-21].
- Michael Stark, Michael Wohlmayr, and Franz Pernkopf. Source–filter-based single-channel speech separation using pitch information. *IEEE Transactions on Audio, Speech, and Language Processing*, 19(2):242–255, 2011.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open Foundation and Fine-Tuned Chat Models, July 2023. arXiv:2307.09288 [cs].
- Efthymios Tzinis, Gordon Wichern, Aswin Subramanian, Paris Smaragdis, and Jonathan Le Roux. Heterogeneous Target Speech Separation. In *Interspeech 2022*, pp. 1–5, September 2022. arXiv:2204.03594 [cs, eess].

- Tuomas Virtanen. Speech recognition using factorial hidden markov models for separation in the feature space. In *Ninth International Conference on Spoken Language Processing*, 2006.
- Tuomas Virtanen. Monaural sound source separation by nonnegative matrix factorization with temporal continuity and sparseness criteria. *IEEE transactions on audio, speech, and language processing*, 15(3):1066–1074, 2007.
- DeLiang Wang and Guy J Brown. *Computational auditory scene analysis: Principles, algorithms, and applications*. Wiley-IEEE press, 2006.
- Xiaofei Wang, Manthan Thakker, Zhuo Chen, Naoyuki Kanda, Sefik Emre Eskimez, Sanyuan Chen, Min Tang, Shujie Liu, Jinyu Li, and Takuya Yoshioka. SpeechX: Neural Codec Language Model as a Versatile Speech Transformer, August 2023. arXiv:2308.06873 [cs, eess].
- Zhong-Qiu Wang, Jonathan Le Roux, and John R. Hershey. Alternative Objective Functions for Deep Clustering. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 686–690, April 2018. ISSN: 2379-190X.
- Mengyue Wu, Heinrich Dinkel, and Kai Yu. Audio Caption: Listen and Tell. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 830–834, May 2019. ISSN: 2379-190X.
- Chenglin Xu, Wei Rao, Eng Siong Chng, and Haizhou Li. SpEx: Multi-Scale Time Domain Speaker Extraction Network. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28: 1370–1384, 2020. ISSN 2329-9304. Conference Name: IEEE/ACM Transactions on Audio, Speech, and Language Processing.
- Takuya Yoshioka, Hakan Erdogan, Zhuo Chen, Xiong Xiao, and Fil Alleva. Recognizing Overlapped Speech in Meetings: A Multichannel Separation Approach Using Neural Networks. In *Interspeech 2018*, pp. 3038–3042. ISCA, September 2018.
- Dong Yu, Morten Kolbaek, Zheng-Hua Tan, and Jesper Jensen. Permutation invariant training of deep models for speaker-independent multi-talker speech separation. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 241–245, New Orleans, LA, March 2017. IEEE. ISBN 978-1-5090-4117-6.
- Dong Zhang, Shimin Li, Xin Zhang, Jun Zhan, Pengyu Wang, Yaqian Zhou, and Xipeng Qiu. SpeechGPT: Empowering Large Language Models with Intrinsic Cross-Modal Conversational Abilities, May 2023. arXiv:2305.11000 [cs].
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging LLM-as-a-judge with MT-Bench and Chatbot Arena, July 2023. arXiv:2306.05685 [cs].
- Katerina Zmolikova, Marc Delcroix, Tsubasa Ochiai, Keisuke Kinoshita, Jan Černocký, and Dong Yu. Neural Target Speech Extraction: An overview. *IEEE Signal Processing Magazine*, 40(3): 8–29, May 2023. ISSN 1558-0792. Conference Name: IEEE Signal Processing Magazine.
- Kateřina Žmolíková, 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*, 13(4): 800–814, August 2019. ISSN 1941-0484. Conference Name: IEEE Journal of Selected Topics in Signal Processing.

## A APPENDIX

### A.1 IMPLEMENTATION DETAILS

**Model Architecture:** The LLM-TSE model incorporates a text cue encoder derived from the LLaMA-2 7B model, a transformer decoder architecture. We generate the text cue embedding using the averaging results of the outputs of the last four self-attention layers. Subsequently, a linear

projection layer is employed to map its dimensions to match the embedding output of the audio cue encoder model. The construction of the audio cue encoder and extractor is built upon an open source of the time-domain SpeakerBeam (TD-SpeakerBeam)<sup>5</sup>. The default model hyperparameters from TD-SpeakerBeam are employed in this process.

**Optimization:** We use the AdamW optimizer for optimization, with an initial learning rate of  $1e-4$ , which has proven effective for various tasks in our preliminary experiments. Our model is trained using ten NVIDIA 3090 GPUs, each with a batch size of 1. For stable training, we employ gradient accumulation, with backpropagation performed every two interactions, culminating in a valid batch size of 40 per iteration. A linear warmup scheduler is used for the first 1000 iteration steps, during which the learning increases from 0 to  $1e-4$  and remains constant. This strategy aims to gradually prepare the model for more complex tasks and improve overall learning stability. Finally, based on our preliminary experiments on the current dataset, we use the gradient normalization with a value of 30. This operation controls the weight update step and prevents gradient explosion.

**LoRA Adaptor:** We adopt the LoRA approach for efficient fine-tuning. The hyperparameters of the LoRA matrix, rank  $r$ , and scaling weight  $\alpha$  are set to 16 and 16. The LoRA dropout is set to 0.05. These LoRA adaptors are applied to the linear projection layers of the query and key calculation in the self-attention layers.

**Generation of Prompts:** We will first write ten question sentences by hand for each task. Then, we used ChatGPT-4-32K to rephrase and generate more questions for each task. The prompt of rephrase is: “Keep it short, limit to 8 words. Feel free to vary sentence structures, but avoid duplications, and synonyms can be replaced. Imitate the tone of a casual conversation, don’t be too rigid. Maintain the existing JSON format when outputting.”

## A.2 LIMITATIONS

While our study introduces a paradigm for speaker extraction, it does not encompass all the acoustical characteristics and contextual cues that can be utilized for this purpose. Indeed, numerous features assist humans in distinguishing auditory objects that we have not fully explored in this research. For instance, pitch, timbre, speech speed rate, and rhythm can provide significant cues for differentiating between speakers (Haykin & Chen, 2005; Mesgarani & Chang, 2012; Popelka et al., 2016). High-level Semantic Information, such as the topic of conversation (e.g., “the person talking about the topic of weather”), can also serve as a powerful identifier. Furthermore, the speaker’s age often influences the characteristics of their speech and can thus be a valuable cue for speaker extraction. These are all potent cues that could be further explored and integrated into our system. Future work could focus on extending the model’s capability to handle these additional attributes, refining the methodology for incorporating such diverse information and evaluating the subsequent improvements in system performance. Moreover, while our proposed system shows promise, there are potential challenges in implementing and testing the system in real-world scenarios, such as in noisy environments or situations with multiple concurrent speakers. Further research is therefore needed to evaluate the system under such conditions and to develop strategies for dealing with these challenges.

In the future, we aim to delve into more open-ended perceptual concepts. Currently, our work is constrained by a reliance on predefined categories, and we cannot handle relatively abstract perceptual descriptions of auditory objects. For example, a description such as “The first speaker’s voice is quite resonant, but after discussing basketball, the voice gradually diminishes” is beyond our current system’s capacity. However, we believe that dealing with more open-ended and advanced problems necessitates a foundational understanding of the basic attributes of auditory objects. This understanding forms the core of our future work. We are committed to exploring how to integrate more open and detailed descriptions of auditory object differences into speech separation scenarios.

<sup>5</sup><https://github.com/BUTSpeechFIT/speakerbeam>