

PROMPT-DRIVEN TARGET SPEECH DIARIZATION

Yidi Jiang¹, Zhengyang Chen², Ruijie Tao^{1*}, Liqun Deng³, Yanmin Qian² and Haizhou Li^{5,4,1}

¹National University of Singapore, Singapore ²Shanghai Jiao Tong University, China
³Huawei Noah's Ark Lab, China ⁴Shenzhen Research Institute of Big Data, Shenzhen, China
⁵School of Data Science, The Chinese University of Hong Kong, Shenzhen, China

ABSTRACT

We introduce a novel task named ‘target speech diarization’, which seeks to determine ‘when target event occurred’ within an audio signal. We devise a neural architecture called Prompt-driven Target Speech Diarization (PTSD), that works with diverse prompts that specify the target speech events of interest. We train and evaluate PTSD using sim2spk, sim3spk and sim4spk datasets, which are derived from the Librispeech. We show that the proposed framework accurately localizes target speech events. Furthermore, our framework exhibits versatility through its impressive performance in three diarization-related tasks: target speaker voice activity detection, overlapped speech detection and gender diarization. In particular, PTSD achieves comparable performance to specialized models across these tasks on both real and simulated data. This work serves as a reference benchmark and provides valuable insights into prompt-driven target speech processing.

Index Terms— Target speech diarization, Prompt-driven, Speaker diarization

1. INTRODUCTION

Humans have the ability to selectively attend to a specific sound source in a complex acoustic environment, that is commonly referred to as the cocktail party effect [1]. Such remarkable auditory attention mechanism allows us to focus our listening effectively [2,3]. Speaker extraction does this when the attended target speaker is known in advance [4–6]. Speaker diarization seeks to demarcate ‘who spoke when’ in a multi-talker speech [7–9]. They serve as the front-end for several speech downstream tasks [10–12]. However, beyond speaker identity [13–16], we are also interested in other semantic types of human speech [17–19], for example, female speech, multi-talker speech mixture, or the speech of keynote speaker who speaks the most in a meeting.

In this paper, we introduce a novel task, termed ‘target speech diarization’, which seeks to determine ‘when target event occurred’ guided by a specific prompt within an audio. From application point of view, this is similar to speech information retrieval where we use a prompt query to retrieve relevant speech segments. From speech processing point of view, this is similar to a speaker extraction task except that the prompt can be specified by a set of speech properties, referred to as semantic attributes, beyond speaker identity. We present a prompt-driven target speech diarization framework that utilizes prompt vectors to provide contextual information as the query. The proposed model architecture is inspired by similar ideas in image segmentation [20–22], audio source separation [23, 24], and speech separation [18].

* Corresponding Author

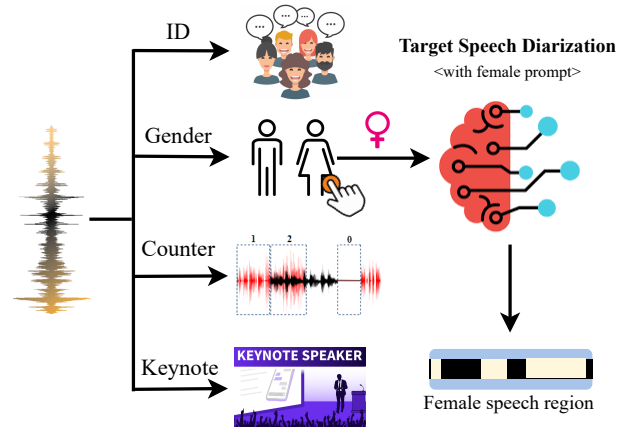


Fig. 1: The illustration of the target speech diarization task. Each semantic attribute (e.g., gender) takes on one or multiple semantic values (e.g., female). The task aims to identify the target event regions given the semantic value information.

Our framework considers four semantic attributes: time-stamped speaker identity, gender, speaker counter (number of speakers at each frame), and keynote speaker (the most talkative speaker). Each attribute can take on one or multiple values, specifying distinct target speech events. We associate each semantic value with its respective prompt vector. By manipulating the combination of semantic attributes, as reflected in the prompts, the proposed model allows us to search over the speech content, thus facilitating a wide range of speech applications.

The contribution of this paper can be summarised as follows:

1. We introduce the innovative task of target speech diarization. Here, we utilize diverse semantic attributes to distinguish different speech events, aligning with human perception and cognitive speech processes.
2. We propose an efficient prompt-driven target speech diarization architecture, effectively detecting target event regions by incorporating heterogeneous prompts query vectors. Meanwhile, we conduct comprehensive experiments to demonstrate the system’s robustness.
3. Our framework extends versatility to target speaker activity detection, overlapped speech detection and gender diarization, each customized to distinct semantic attribute. Notably, we also conduct comparative analysis with specialized models on both real and simulated datasets across these tasks.

2. TARGET SPEECH DIARIZATION

2.1. Task formulation

To formulate our task, we first introduce two concepts, the semantic attribute and semantic value. The semantic attribute represents the criterion of demarcating speech segments. Each semantic attribute takes on one or multiple semantic values associated with specific events. For examples, in speaker diarization task, speaker identity is the semantic attribute. The specific speaker ID is semantic value and his/her speaking activity is the aligned speech event.

As we mentioned in Section 1, an audio can be characterized by various semantic attributes. In our proposed target speech diarization task, the system will simultaneously take audio and semantic value information as inputs and output target event regions related to this semantic value. For example, if the semantic value is female from the semantic attribute gender, the system should output the entire female-speaking regions. To demonstrate the feasibility of the task, we consider four semantic attributes in our work, denoted as \mathcal{T} , \mathcal{G} , \mathcal{N} and \mathcal{K} :

- \mathcal{T} - **Timestamped speaker identity:**

In this attribute, the attribute values consist of timestamps that point to individual speakers. We use the brief timestamp-based description “the person who spoke at the particular time” to specify the speaker identity.

Compared with traditional approaches which always rely on the pre-enrolled speaker embeddings, our system is flexible and user-friendly for real-world applications. For instance, when we seek to locate all speech segments for a person of interest, we can simply scan the audio to find a timestamp when he/she is the only one talking.

- \mathcal{G} - **Gender:**

The gender attribute is more straightforward and contains two values, female and male, which can guide the system to output the gender-specific event regions.

- \mathcal{N} - **Speaker counter:**

This attribute identifies the number of concurrent speakers at each frame and contains three event values: non-speech, single-speaker speech, and overlapped speech.

- \mathcal{K} - **Keynote speaker:**

This attribute focuses on identifying the keynote speaker. It contains one event value to represent the person who talks most.

Identifying the keynote speaker is crucial for real-world applications. By leveraging both keynote and speaker counter prompts, user can differentiate between speech segments belonging to the keynote speaker and others.

2.2. Proposed framework: Prompt-driven Target Speech Diarization (PTSD)

To solve the task we formulated in the previous section, we proposed a framework called Prompt-driven Target Speech Diarization (PTSD). In this framework, we modeled each semantic value information as a prompt vector $p \in \mathbb{R}^{1 \times D}$. Such a prompt-style framework [20, 25] can be the basis for constructing a versatile and flexible system.

In this paper, we denoted the audio input as X and formulated the target event region aligned with each semantic value as a binary sequence $y \in \{0, 1\}^{1 \times T}$, where 1 represents the existence of target event and 0 represents the absence. T is the number of audio frame.

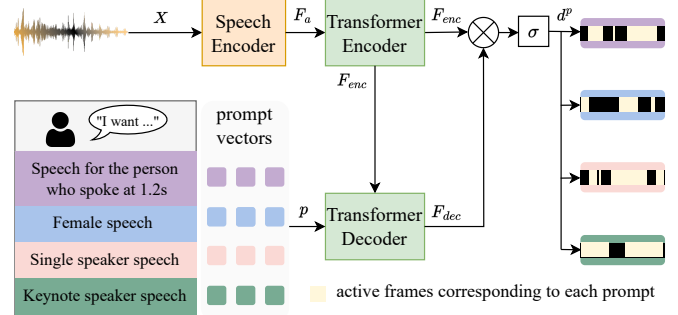


Fig. 2: The overview of our prompt-driven target speech diarization framework. It takes audio and prompt vectors according to the user intention as inputs, and outputs target event regions aligned with the prompt vectors. \otimes and σ represent the dot product and sigmoid operation, respectively.

As depicted in Figure 2, PTSD comprises a speech encoder and a transformer encoder-decoder. The speech encoder first maps the input audio X to a feature sequence $F_a \in \mathbb{R}^{T \times D}$. Then, the transformer encoder-decoder takes F_a and prompt vector p as input and outputs the prediction for target event related to p . Specifically, the transformer encoder takes F_a as input and outputs the frame-level speech representation $F_{enc} \in \mathbb{R}^{T \times D}$. The transformer decoder takes prompt vector p and F_{enc} as inputs, and output $F_{dec} \in \mathbb{R}^{1 \times D}$. Finally, we performed a dot product operation between the decoder output F_{dec} and the encoder output F_{enc} and applied a sigmoid operation to get the prediction sequence $d^p \in (0, 1)^{1 \times T}$. The value of d^p denotes the target event occurrence probability at each frame. Notably, our framework can support one or multiple prompt vectors at the same and output their associated target events regions accordingly.

Speech encoder. In our framework, we employed a pre-trained WavLM encoder [26] as the speech encoder to obtain frame-level representations F_a . With consideration for the trade-off between computational efficiency and speech information, we utilized the first three layers WavLM encoder, freezing them during our training process. The WavLM encoder [26] was designed to learn universal speech representations from vast amounts of unlabeled speech data, ensuring the universality and robustness of the frame-level audio representations.

Prompt query vectors. Each target event value is associated with its own prompt query vector. By switching between different prompts, our framework can accurately detect distinct event regions corresponding to each prompt query vector p .

For prompts belonging to the timestamped speaker attribute, we extracted a single vector from the temporal axis of the frame-level representation F_a using the provided particular frame index for the timestamped speaker. This extracted vector serves as the prompt query vector of this timestamped speaker. For prompts related to gender attribute, we used two learnable embeddings to present male and female semantic values separately during the training stage. Similarly, for other attributes such as \mathcal{N} and \mathcal{K} , we also employed learnable embeddings to provide information for each semantic value.

Transformer encoder-decoder. We employed the transformer encoder-decoder architecture as introduced in [27]. Leveraging the power of self-attention and cross-attention mechanisms, the transformer encoder-decoder excels in capturing intricate temporal patterns in the audio data and aligning them with the relevant prompts. This synergy enables our model to precisely identify and diarize tar-

get event regions, making it a robust and adaptable solution for our task.

Loss function. The learning targets of our framework are frame-wise binary ground truth labels $y \in \{0, 1\}^{N \times T}$ of N target events. For each target event, we utilized binary cross-entropy loss to train our model, as defined in Equation 1. d_t and y_t represent the predicted and ground truth labels of a specific target event for the t^{th} audio frame, where $t \in [1, T]$. The loss function is designed to minimize the difference between predicted and ground truth labels, encouraging our model to accurately detect target event activities.

$$\mathcal{L} = -\frac{1}{T} \sum_{t=1}^T (y_t \cdot \log d_t) + (1 - y_t) \cdot \log(1 - d_t) \quad (1)$$

3. EXPERIMENTS

In this section, we detailed the datasets, evaluation metrics, and experimental setup used to evaluate the proposed PTSD framework.

3.1. Dataset

Given that real-world speech datasets cannot provide all the required groundtruth labels according to the semantic attributes introduced in Section 2.1, we followed the recipe¹ proposed in [28] to simulate 2-, 3-, and 4-speaker datasets from Librispeech [29]. To create datasets that closely resemble real-world conversations, we utilized conversation statistics from the DIHARD II development set [30] to generate 1000 hours of audio for each sim2spk, sim3spk, and sim4spk dataset.

3.2. Evaluation metric

For target speech diarization evaluation, we primarily employed three metrics: accurate precision (AP), area under the receiver operating characteristic (AUC), and equal error rate (EER) based on the implementation from sklearn package.

3.3. Implementation details

The proposed PTSD framework was implemented using PyTorch and optimized with the Adam optimizer. We set the initial learning rate to 10^{-4} and decrease it by 5% for each epoch. The dimension D of audio feature F_a and prompt query vector p were both set to 256. For both transformer encoder and decoder structure, 4-layer transformer with 8 attention heads was applied. To ensure the robustness of our system, we conducted experiments using inputs of diverse lengths, spanning from 20 to 60 seconds during the training phase. For validation purposes, all inputs were segmented into 40-second chunks for simplicity.

4. RESULTS AND ANALYSIS

In this section, we presented a comprehensive demonstration of our framework’s performance across heterogeneous prompts to show the feasibility of our proposed target speech diarization task. Furthermore, our framework’s applicability can be extended to target speaker activity detection, concurrent speaker counting and gender diarization tasks, each aligned with specific semantic attribute. Moreover, we conducted a comparative analysis with the specialists model to evaluate the effectiveness of our approach.

4.1. Overall analysis of PTSD

In this section, we conducted the training phase of our framework using the sim2spk, sim3spk, and sim4spk datasets. The performances are depicted in Table 1, showcasing the AP, AUC, EER results for the prompts from four different semantic attributes. Notably, in the case of sim4spk, our model can support ten prompts vectors to specify ten

Table 1: Overall analysis of PTSD system on the sim2spk, sim3spk and sim4spk datasets across diverse semantic attributes.

Dataset	Attribute	AP (%) \uparrow	AUC (%) \uparrow	EER (%) \downarrow
sim2spk	\mathcal{T}	99.90	99.91	1.18
	\mathcal{G}	95.65	96.38	5.98
	\mathcal{N}	99.84	99.92	1.47
	\mathcal{K}	99.91	99.84	1.68
sim3spk	\mathcal{T}	99.40	99.65	2.80
	\mathcal{G}	95.70	96.46	5.93
	\mathcal{N}	99.56	99.77	2.41
	\mathcal{K}	99.32	99.21	4.81
sim4spk	\mathcal{T}	98.88	99.57	3.29
	\mathcal{G}	96.70	97.19	6.20
	\mathcal{N}	99.53	99.76	2.51
	\mathcal{K}	98.75	98.88	5.85

target events regions simultaneously, comprising four timestamped speakers (\mathcal{T}), two under gender (\mathcal{G}) (related to female and male), three under speaker counter attribute (\mathcal{N}) (related to non-speech, single speaker speech, overlapped speech) and one under keynote speaker attribute (\mathcal{K}).

Impressively, all the AP and AUC values surpass 95%, and EER values are under 7%. These results demonstrate that our PTSD framework can accurately identify the desired event regions guided by provided prompt query vectors. In this way, we can switch the prompt query vector according to the user intention, which is flexible and powerful.

4.2. PTSD with specific semantic attribute

The results from the previous section indicate that we can utilize a single PTSD model by switching between different prompts to detect various speech events, and it has demonstrated commendable overall performance. To further evaluate PTSD, in this section, we compared PTSD with specialized models for different sub-tasks with specific semantic attribute.

4.2.1. PTSD (\mathcal{T}) v.s. TS-VAD

Our framework also has the potential to perform the target speaker voice activity detection (TS-VAD) task when furnished with timestamp for each speaker, denoted as PTSD (\mathcal{T}). In the primitive TS-VAD paradigm [9], an i-vector for the target speaker was provided to help the system to detect the speaking activity.

The key difference between TS-VAD and PTSD lies in the methods used to provide the speaker prior information: timestamped speaker prompt vectors in PTSD (\mathcal{T}) and speaker embeddings in TS-VAD.

To ensure a fair comparison, we made modifications to the original TS-VAD model and named it as ‘mod. TS-VAD’. Specifically, we replaced the Bidirectional Long Short-Term Memory (BLSTM) and MFCC used in the original TS-VAD with a four-layer transformer encoder and a three-layer WavLM encoder, respectively. Additionally, we employed the ECAPA-TDNN² trained on a combination of VoxCeleb, CnCeleb [31] and Alimeeting training sets, to extract accurate speaker embeddings for TS-VAD.

In PTSD, the timestamped speaker prompt vector corresponds to a temporal duration of 0.04 seconds (given that the length of the WavLM feature is 25 for one second) within the audio sequence. To ensure an objective comparison, we selected clean speech segments

¹https://github.com/BUTSpeechFIT/EEND_dataprep/

²<https://github.com/TaoRuijie/ECAPA-TDNN>

of length 1s, 2s, and 3s and employed them in the ‘mod. TS-VAD’ system to extract speaker embedding. We provided these speaker priors for each input chunk to conduct the comparison experiments on Alimeeting and sim4spk datasets. Besides, since TS-VAD system was first introduced in a speaker diarization task [32], we used the diarization error rate (DER) as the evaluation metric for comparison.

The results, as shown in Table 2, highlight that PTSD (\mathcal{T}) outperforms all three versions of TS-VAD on sim4spk dataset. Furthermore, PTSD also performs better than TS-VAD with 1s enrollment speech and achieves comparable results with 2s and 3s enrollment speech on Alimeeting dataset. This difference can be attributed to the ECAPA-TDNN encoder, which has been fine-tuned on the Alimeeting dataset. In contrast, our method does not rely on a specific speaker encoder, offering enhanced flexibility and convenience in this regard.

Table 2: The performance comparison between PTSD (\mathcal{T}) and TS-VAD on Alimeeting and sim4spk datasets.

Dataset	Method	DER (%) \downarrow
Alimeeting	mod. TS-VAD (with 1s ref)	12.63
	mod. TS-VAD (with 2s ref)	10.22
	mod. TS-VAD (with 3s ref)	8.80
	PTSD (\mathcal{T})	11.40
sim4spk	mod. TS-VAD (with 1s ref)	12.63
	mod. TS-VAD (with 2s ref)	8.81
	mod. TS-VAD (with 3s ref)	7.03
	PTSD (\mathcal{T})	5.58

It should be noting that the TS-VAD system cannot perform speaker diarization task independently. The authors in TS-VAD [9] leveraged a pre-trained diarization system to provide clean enrollment speech for each speaker, and then TS-VAD system was applied to detect the activity of each speaker. As an initial attempt, our experiments have demonstrated the feasibility of using timestamps as speaker enrollment priors to perform target speaker activity detection. Similarly, we can also extend our PTSD system in the same way as [9] to complete the speaker diarization task. We will continue to explore the timestamp-based speaker diarization system in the future.

4.2.2. PTSD (\mathcal{N}) v.s. OSD

PTSD can also be functioned as a three-class speaker counter, capable of estimating the number of concurrent speakers at each frame when we provide prompt vectors for non-speech, single speaker speech and overlapped speech simultaneously. We denote this mode as PTSD (\mathcal{N}).

We evaluated the performance of PTSD in the overlapped speech detection (OSD) task on DIHARD II evaluation set. We used the overlapped speech prompt vector, which was initially trained on the sim4spk training set and further fine-tuned on the DIHARD II development set. Table 3 presents the comparison results between PTSD (\mathcal{N}) and previous two OSD models proposed by Jung et al [33] and Bullock et al [34] respectively. PTSD (\mathcal{N}) achieves significantly better precision at 68.93% and recall at 48.18% compared to the specialized overlapped speech detection models on DIHARD II evaluation set.

4.2.3. Gender diarization: PTSD (\mathcal{G})

When we provide both female and male prompts, PTSD (\mathcal{G}) can perform for gender diarization for the first attempt to answer the question: “which gender appeared when”. We implemented the ‘baseline1’ using WavLM speech encoder and ECAPA-TDNN encoder

Table 3: Overlapped speech detection comparison results on DIHARD II evaluation set.

Method	Precision (%) \uparrow	Recall (%) \uparrow
Bullock et al. [34]	64.50	26.70
Jung et al. [33]	66.48	32.22
PTSD (\mathcal{N})	68.93	48.18

followed by fully connection layer as frame-wise binary classification. The second baseline we implemented is using WavLM speech encoder and transformer encoder with fully connection layer, denoted as ‘baseline2’. As shown in Table 4, PTSD (\mathcal{G}) can get better performance than two baselines. At the same time, our framework can obtain more flexible prompt-driven outputs with transformer decoder structure.

Table 4: Gender diarization on sim4spk datasets.

Method	AP (%) \uparrow	AUC (%) \uparrow	EER (%) \downarrow	DER (%) \downarrow
baseline1	96.85	97.85	5.50	8.45
baseline2	97.31	97.57	6.34	9.04
PTSD (\mathcal{G})	98.17	98.39	5.13	7.75

4.3. Discussion and future work

We believe that the attention mechanism plays a crucial role in making our framework operate effectively. Taking PTSD (\mathcal{T}) as an example, the attention structure combines information from nearby audio frames into the specific frame, thereby enriching the speaker-related information.

In the future, our research will progress toward integrating natural language commands and supporting various prompt forms. This expansion aims to improve the input query’s effectiveness by adopting a multi-modal approach, which should enhance the system’s adaptability and versatility in real-world applications.

5. CONCLUSION

In this paper, we have introduced an innovative task called “target speech diarization”, aimed at distinguishing diverse speech events from various perspectives. This task mimics how humans naturally engage with audio in their daily lives. Additionally, we have proposed a framework, Prompt-driven Target Speech Diarization (PTSD), to replicate the multi-dimensional auditory comprehension process observed in humans. We have developed specific prompts for each target event, allowing us to switch between different functional modes. Our model’s performance across various semantic attributes and the subsequent comparison with specialized models have demonstrated the superior performance and flexibility of our approach. Our study provides new insights for diarization-related tasks. More practical application scenarios can be designed based on this direction.

6. ACKNOWLEDGEMENTS

This work is supported by Huawei Noah’s Ark Lab, Human Robot Collaborative AI under its AME Programmatic Funding Scheme (Project No. A18A2b0046), the National Natural Science Foundation of China (Grant No. 62271432), Shenzhen Science and Technology Research Fund (Fundamental Research Key Project Grant No. JCYJ20220818103001002), the Internal Project Fund from Shenzhen Research Institute of Big Data under Grant No. T00120220002, and the Advanced Research and Technology Innovation Centre (ARTIC), the National University of Singapore under Grant (project number: A0005947-21-00, project reference: ECT-RP2).

7. REFERENCES

- [1] E Colin Cherry, “Some experiments on the recognition of speech, with one and with two ears,” *The Journal of the acoustical society of America*, vol. 25, no. 5, pp. 975–979, 1953.
- [2] Jonathan B Fritz, Mounya Elhilali, Stephen V David, and Shihab A Shamma, “Auditory attention—focusing the searchlight on sound,” *Current opinion in neurobiology*, vol. 17, no. 4, pp. 437–455, 2007.
- [3] Nima Mesgarani and Edward F Chang, “Selective cortical representation of attended speaker in multi-talker speech perception,” *Nature*, pp. 233–236, 2012.
- [4] Junjie Li, Meng Ge, Zexu Pan, Rui Cao, Longbiao Wang, Jianwu Dang, and Shiliang Zhang, “Rethinking the Visual Cues in Audio-Visual Speaker Extraction,” in *Proc. INTERSPEECH 2023*, 2023, pp. 3754–3758.
- [5] 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*, vol. 28, pp. 1370–1384, 2020.
- [6] Meng Ge, Chenglin Xu, Longbiao Wang, Eng Siong Chng, Jianwu Dang, and Haizhou Li, “SpEx+: A complete time domain speaker extraction network,” *arXiv preprint arXiv:2005.04686*, 2020.
- [7] Zhengyang Chen, Bing Han, Shuai Wang, and Yanmin Qian, “Attention-based encoder-decoder network for end-to-end neural speaker diarization with target speaker attractor,” *INTER_SPEECH*, 2023.
- [8] Ming Cheng, Weiqing Wang, Yucong Zhang, Xiaoyi Qin, and Ming Li, “Target-speaker voice activity detection via sequence-to-sequence prediction,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [9] Ivan Medennikov, Maxim Korenevsky, Tatiana Prisyach, Yuri Khokhlov, Mariya Korenevskaya, Ivan Sorokin, Tatiana Timofeeva, Anton Mitrofanov, Andrei Andrusenko, Ivan Podluzhny, et al., “Target-speaker voice activity detection: a novel approach for multi-speaker diarization in a dinner party scenario,” *INTER_SPEECH*, 2020.
- [10] Xiaoxue Gao, Chitralkha Gupta, and Haizhou Li, “Automatic lyrics transcription of polyphonic music with lyrics-chord multi-task learning,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 2280–2294, 2022.
- [11] Hexin Liu, Leibny Paola Garcia, Xiangyu Zhang, Andy W. H. Khong, and Sanjeev Khudanpur, “Enhancing code-switching speech recognition with interactive language biases,” *arXiv preprint arXiv:2309.16953*, 2023.
- [12] Yidi Jiang, Bidisha Sharma, Maulik Madhavi, and Haizhou Li, “Knowledge distillation from bert transformer to speech transformer for intent classification,” *arXiv preprint arXiv:2108.02598*, 2021.
- [13] Brecht Desplanques, Jenthe Thienpondt, and Kris Demuyne, “Ecapadnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification,” *arXiv preprint arXiv:2005.07143*, 2020.
- [14] Yidi Jiang, Ruijie Tao, Zexu Pan, and Haizhou Li, “Target active speaker detection with audio-visual cues,” *arXiv preprint arXiv:2305.12831*, 2023.
- [15] Tianchi Liu, Kong Aik Lee, Qiongqiong Wang, and Haizhou Li, “Disentangling voice and content with self-supervision for speaker recognition,” in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [16] Tianchi Liu, Kong Aik Lee, Qiongqiong Wang, and Haizhou Li, “Golden gemini is all you need: Finding the sweet spots for speaker verification,” *arXiv preprint arXiv:2312.03620*, 2023.
- [17] Yingzhi Wang, Mirco Ravanelli, Alaa Nfissi, and Alya Yacoubi, “Speech emotion diarization: Which emotion appears when?,” *arXiv preprint arXiv:2306.12991*, 2023.
- [18] Efthymios Tzinis, Gordon Wichern, Aswin Subramanian, Paris Smaragdis, and Jonathan Le Roux, “Heterogeneous target speech separation,” *INTER_SPEECH*, 2022.
- [19] Martin Lebourdais, Marie Tahon, Antoine Laurent, and Sylvain Meignier, “Overlapped speech and gender detection with wavlm pre-trained features,” *arXiv preprint arXiv:2209.04167*, 2022.
- [20] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al., “Segment anything,” *arXiv preprint arXiv:2304.02643*, 2023.
- [21] Xinyin Ma, Gongfan Fang, and Xinchao Wang, “Deepcache: Accelerating diffusion models for free,” *arXiv preprint arXiv:2312.00858*, 2023.
- [22] Gongfan Fang, Xinyin Ma, Mingli Song, Michael Bi Mi, and Xinchao Wang, “DepGraph: Towards Any Structural Pruning,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [23] 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.
- [24] Ying Hu, Yadong Chen, Wenzhong Yang, Liang He, and Hao Huang, “Hierarchic temporal convolutional network with cross-domain encoder for music source separation,” *IEEE Signal Processing Letters*, vol. 29, pp. 1517–1521, 2022.
- [25] Brian Lester, Rami Al-Rfou, and Noah Constant, “The power of scale for parameter-efficient prompt tuning,” *arXiv preprint arXiv:2104.08691*, 2021.
- [26] Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, 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.
- [27] 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.
- [28] Federico Landini, Mireia Diez, Alicia Lozano-Diez, and Lukáš Burget, “Multi-speaker and wide-band simulated conversations as training data for end-to-end neural diarization,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [29] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [30] Neville Ryant, Kenneth Church, Christopher Cieri, Alejandrina Cristia, Jun Du, Sriram Ganapathy, and Mark Liberman, “The second dihard diarization challenge: Dataset, task, and baselines,” *arXiv preprint arXiv:1906.07839*, 2019.
- [31] Yue Fan, JW Kang, LT Li, KC Li, HL Chen, ST Cheng, PY Zhang, ZY Zhou, YQ Cai, and Dong Wang, “Cn-celeb: a challenging chinese speaker recognition dataset,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7604–7608.
- [32] Shota Horiguchi, Yusuke Fujita, Shinji Watanabe, Yawen Xue, and Kenji Nagamatsu, “End-to-end speaker diarization for an unknown number of speakers with encoder-decoder based attractors,” *arXiv preprint arXiv:2005.09921*, 2020.
- [33] Jee-weon Jung, Hee-Soo Heo, Youngki Kwon, Joon Son Chung, and Bong-Jin Lee, “Three-class overlapped speech detection using a convolutional recurrent neural network,” *arXiv preprint arXiv:2104.02878*, 2021.
- [34] Latané Bullock, Hervé Bredin, and Leibny Paola Garcia-Perera, “Overlap-aware diarization: Resegmentation using neural end-to-end overlapped speech detection,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7114–7118.