

TellWhisper: Tell Whisper Who Speaks When

Yifan Hu^{1,2}, Peiji Yang², Zhisheng Wang², Yicheng Zhong², Rui Liu^{1*}

¹ Inner Mongolia University, Hohhot, China

² Tencent Technology Co.Ltd, Shenzhen, Guangdong, China

22309013@mail.imu.edu.cn, imucslr@imu.edu.cn,
{peijiyang, plorywang, ajaxzhong}@tencent.com

Abstract

Multi-speaker automatic speech recognition (MASR) aims to predict “*who spoke when and what*” from multi-speaker speech, a key technology for multi-party dialogue understanding. However, most existing approaches decouple temporal modeling and speaker modeling when addressing “*when*” and “*who*”: some inject speaker cues before encoding (e.g., speaker masking), which can cause irreversible information loss; others fuse identity by mixing speaker posteriors after encoding, which may entangle acoustic content with speaker identity. This separation is brittle under rapid turn-taking and overlapping speech, often leading to degraded performance. To address these limitations, we propose **TellWhisper**, a unified framework that jointly models speaker identity and temporal within the speech encoder. Specifically, we design **TS-RoPE**, a time-speaker rotary positional encoding: time coordinates are derived from frame indices, while speaker coordinates are derived from speaker activity and pause cues. By applying region-specific rotation angles, the model explicitly captures per-speaker continuity, speaker-turn transitions, and state dynamics, enabling the attention mechanism to simultaneously attend to “*when*” and “*who*”. Moreover, to estimate frame-level speaker activity, we develop **Hyper-SD**, which casts speaker classification in hyperbolic space to enhance inter-class separation and refine speaker-activity estimates. Extensive experiments demonstrate the effectiveness of the proposed approach. The project webpage is available at <https://walker-hyf.github.io/TellWhisper>.

1 Introduction

Multi-speaker Automatic Speech Recognition (MASR) aims to predict who speaks what content and at what time in speech containing interactions among multiple speakers (Polok et al., 2025; Yin

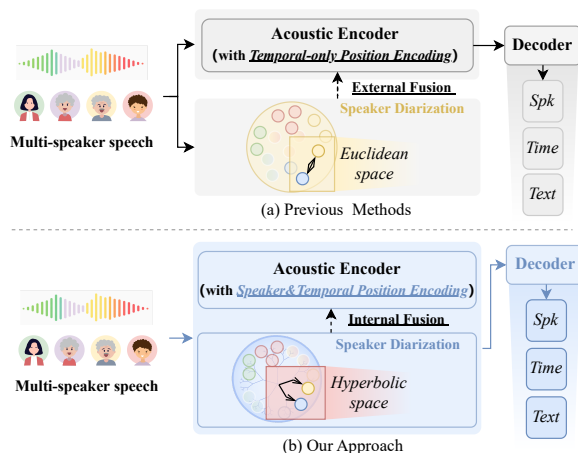


Figure 1: (a) Prior methods model temporal structure and speaker information separately. (b) Our approach uses a unified positional encoding to capture both temporal and speaker dynamics.

et al., 2025). It is a complex task that jointly integrates speaker diarization (SD) (Bredin et al., 2020) and automatic speech recognition (ASR) (Cao et al., 2012). With the development of speech intelligence and conversational systems, MASR plays an increasingly critical role in meeting and interview transcription (Vinnikov et al., 2024), multi-user human-computer interaction (Shin et al., 2025), and the construction of data for spoken dialogue speech foundation models (Ju et al., 2025; Xie et al., 2025; Hao et al., 2025; Weihua et al., 2026). Consequently, developing efficient and robust MASR models is of practical importance.

While current ASR models (Yao et al., 2023; Xu et al., 2025) excel at recognizing linguistic content, their performance often degrades markedly in multi-party dialogues with rapid speaker-turn taking, largely because the critical cues of “*who*” and “*when*” remain insufficiently modeled. In MASR, traditional solutions typically fuse SD and ASR outputs in parallel: the former predicts speaker identities and timestamps, the latter predicts content

*Corresponding author.

and timestamps, and the two streams are aligned by timestamps (Yamasaki et al., 2023). However, accurate timestamp alignment is challenging, especially under overlapping speech, and this pipeline often results in incorrect speaker assignment. Recent works seek to unify SD and ASR, yet most approaches remain fundamentally factorized, modeling temporal structure and speaker identity separately and aggregating speaker cues with acoustic representations *outside* the encoder. As shown in Fig. 1, they use absolute positional encoding for time modeling and adopt three common speaker strategies: (1) (Polok et al., 2025) masks non-target regions before encoding using SD labels, to preserve temporal, blank inputs are still decoded, which can trigger hallucinations. (2) (Kang et al., 2025; He et al., 2025) attempts to isolate the target speaker, but requires extra speaker prompts (Ma et al., 2024; Guo et al., 2024) or fixed number of separated individuals (Zhao and Ma, 2023), and struggles in overlapping regions. (3) Other methods (Park et al., 2024; Medennikov et al., 2025) add predefined speaker sinusoidal kernels weighted by posteriors to encoder states, such linear mixing entangles semantics with speaker cues and complicates decoding. Therefore, how can we model temporal and speaker jointly *within* the encoder in a more seamless way?

To overcome factorized modeling, we propose **TellWhisper** (Fig. 1, lower). The model injects temporal and speaker information into the ASR encoder via positional encoding. Specifically, we design **TS-RoPE**, a time-speaker-aware rotary positional encoding, and apply it to encoder self-attention to modulate Query-Key dot products through controllable rotation-angle differences. We partition the Query/Key channels into temporal subspaces indexed by absolute frame time and speaker subspaces derived from per-frame activity to capture speaker-state dynamics (e.g., sustained speech and pauses). We also allocate disjoint channel regions to different speakers to avoid inter-speaker interference. To obtain more reliable frame-level activity, we further propose **Hyper-SD**, which replaces Euclidean linear scoring with a hyperbolic “feature-prototype distance” (red box in Fig. 1). Negative curvature induces exponential volume growth, so small shifts yield larger distance changes, improving separability among timbrally similar speakers and stabilizing speaker posteriors.

In summary, the main contributions of this paper are as follows:

- We propose TellWhisper, a novel multi-speaker ASR model that introduces TS-RoPE, a time-speaker-aware rotary positional encoding, into the speech encoder to naturally integrate temporal and speaker activity.
- To obtain reliable frame-level speaker activity, we develop Hyper-SD, a hyperbolic-space speaker diarization model that estimates speaker activity via “feature-prototype distances.”
- We conduct extensive experiments that demonstrate the effectiveness of TS-RoPE for time-speaker integration and show that Hyper-SD provides reliable speaker-activity estimates.

2 Related Works

2.1 Rotational Position Encoding

Traditional absolute positional encoding (PE) injects fixed position-dependent vectors into semantic representations (Vaswani et al., 2017), requiring a predefined maximum length and failing to explicitly model relative positions. In contrast, Rotary Positional Embedding (RoPE) (Su et al., 2024) rotates Query and Key vectors so attention depends on relative angles, preserves norms, and supports long context. Beyond large language models (Bai et al., 2023; Touvron et al., 2023), RoPE also applies to speech tasks such as ASR (Zhang et al., 2025) and speech enhancement (Chen and Wang, 2024), where frame features rotate by time to encode dynamics. In vision, RoPE extends to multi-dimensional variants that encode multiple axes (Lu et al., 2024). More recently, multi-dimensional RoPE (Yang et al., 2025) unifies positional encoding across modalities by partitioning channels into semantic subspaces (e.g., *width* and *height*) and encoding factors independently within shared attention. Motivated by these advances, we target MASR, which requires joint temporal and speaker modeling. Instead of encoding time alone, we split channels into temporal and speaker subspaces: the temporal subspace uses standard time rotation, while the speaker subspace is modulated by speaker activity.

2.2 Hyperbolic Representation Learning and Classification

Conventional classifiers (Bredin et al., 2020) typically use a linear head in Euclidean space, but

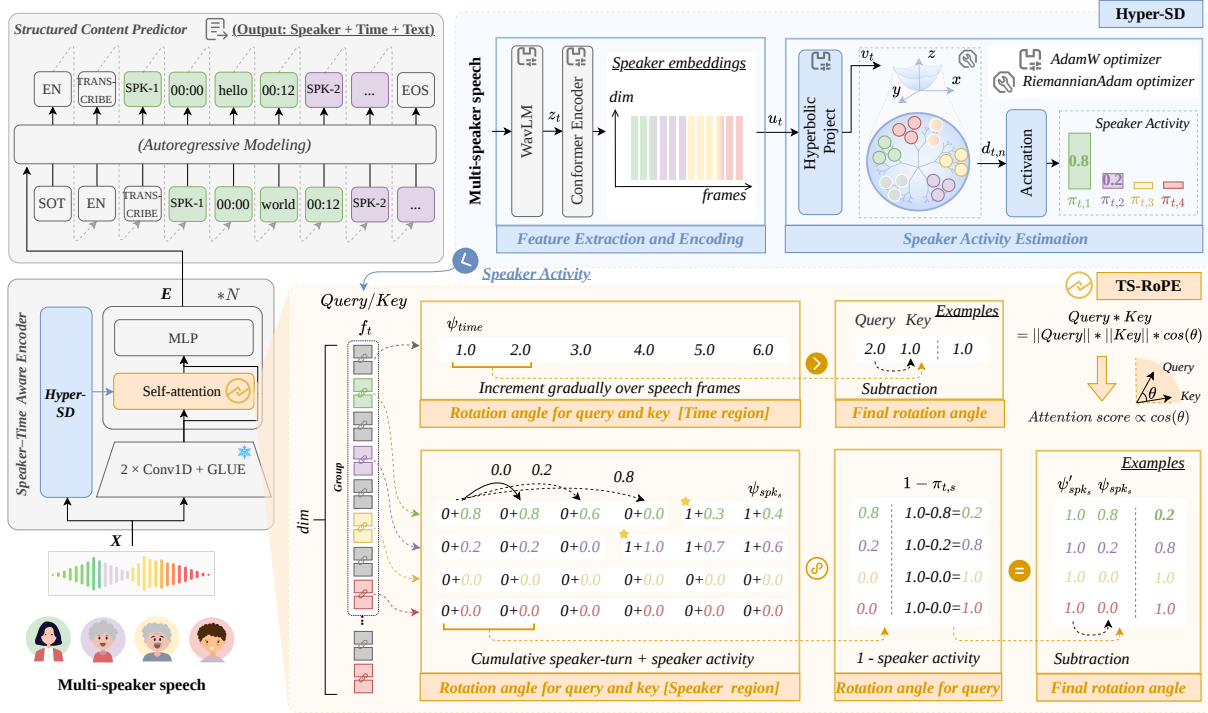


Figure 2: Overall architecture of the TellWhisper model. For multi-speaker speech, the Speaker-Time Aware Encoder encodes the input with convolutional layers and uses Hyper-SD to estimate frame-level speaker activity. Guided by TS-RoPE, self-attention jointly models temporal and speaker dynamics, and the Structured Content Predictor outputs speaker, time, and text. In particular, TS-RoPE builds separate temporal and speaker coordinates and encodes them into disjoint Query/Key subspaces, strengthening attention for aligning “when” and “who” cues.

Euclidean geometry’s flatness and polynomial volume growth make it hard to form large inter-class margins for highly similar distributions, leading to poor discrimination (Xu et al., 2023). Hyperbolic space, with negative curvature and exponential volume expansion, amplifies distance contrasts and enlarges margins (Ganea et al., 2018), which benefits speaker diarization where similar timbres produce confusable embeddings. Hyperbolic embeddings also capture hierarchical structure with low distortion (Pal et al., 2024) and use geometric cues (e.g., radius) to reflect a continuum from ambiguous to separable events (Petermann and Kim, 2024). However, SD requires explicit frame-level discrete labels: non-speaking (noise/silence) should be grouped, while different overlap patterns (e.g., “spk-A & spk-B” vs. “spk-B & spk-C”) must remain separable (Bredin et al., 2020). If ambiguous segments collapse near the origin, separability across overlap types degrades. Accordingly, we assign distinct labels to non-speaking segments, each single speaker, and each overlap combination, and enforce supervision that pushes features and prototypes toward well-separated boundary regions. Finally, we compute frame-level speaker activity

from feature–prototype distances.

3 Task Definition

In multi-speaker automatic speech recognition, the input is a multi-speaker speech signal represented as a frame-level acoustic feature sequence $X = \{x_t\}_{t=1}^T$, where $x_t \in \mathbb{R}^D$ is the feature vector of frame t and T is the sequence length. The signal may contain overlap, rapid speaker transitions, and silence (non-speaking). The MASR model aims to infer structured outputs (speaker identities, timestamps, and transcribed text), formulated as

$$Y = \{(\text{spk}_s, [\tau_{s,j}^{\text{start}}, \mathbf{y}_{s,j}, \tau_{s,j}^{\text{end}}])_{j=1}^J\}_{s=1}^S \quad (1)$$

where spk_s denotes the speaker label, $\tau_{s,j}^{\text{start}}$ and $\tau_{s,j}^{\text{end}}$ are the segment boundaries for the j -th turn of speaker s , $\mathbf{y}_{s,j}$ is the associated text sequence, J is the number of speaker-turn segments of spk_s in X , and S is the number of speakers in X .

4 Proposed Approach: TellWhisper

As shown in Fig. 2, we present the overall architecture of **TellWhisper**. We first describe how Hyper-SD estimates speaker activity, and then introduce

the TS-RoPE-based time-speaker-aware encoder and the structured content predictor.

4.1 Frame-level Speaker Activity Estimator

As shown in the upper-right of Fig. 2, Hyper-SD consists of two stages: (1) it learns speech representations from multiple WavLM (Chen et al., 2022) layers and uses a Conformer encoder to inject global context into frame-level features; (2) a hyperbolic classifier maps Euclidean features into hyperbolic space and estimates speaker activity via feature-prototype distances.

4.1.1 Speech Feature Extraction and Encoding

Given multi-speaker speech X , we use WavLM to extract multi-layer frame-level representations $\mathbf{h}_t^{(l)}$. A learnable weighted-sum aggregation fuses these features into a compact frame representation:

$$\mathbf{z}_t = \sum_{l=1}^L \alpha_l \mathbf{h}_t^{(l)} \quad (2)$$

where, l is the layer index, t is the frame index, and α_l denotes the layer weight.

The aggregated features are then fed into a Conformer to model contextual dependencies:

$$\mathbf{u}_{1:T} = \text{Conformer}(\mathbf{z}_{1:T}) \quad (3)$$

where T is the number of frames. The Conformer integrates long-range context and local acoustic patterns to produce context-aware frame representations for speaker activity estimation.

4.1.2 Prototype-Based Speaker Activity Estimation

Speaker activity estimation is performed in hyperbolic space. Specifically, we first apply a linear transformation and norm clipping to the Euclidean feature \mathbf{u}_t :

$$\begin{aligned} \mathbf{v}_t &= \mathbf{W}\mathbf{u}_t + \mathbf{b} \in \mathbb{R}^I, \\ \mathbf{v}_t &= \mathbf{v}_t \cdot \min\left(1, \frac{r}{\|\mathbf{v}_t\|_2 + \epsilon}\right) \end{aligned} \quad (4)$$

Here, I denotes the hyperbolic embedding dimension, r controls the clipping radius, \mathbf{W} and \mathbf{b} are the weight matrix and bias, and ϵ is a small constant.

A Poincaré ball (Ungar, 2001) \mathbb{B}_c with curvature c serves as the underlying hyperbolic space. The clipped features are mapped to \mathbb{B}_c via the exponential map at the origin and then projected to

remain inside the ball for numerical stability. We assign a learnable hyperbolic prototype $\mathbf{p}_n \in \mathbb{B}_c$ to each speaker-combination¹ class $n \in \mathcal{N}$, where \mathcal{N} is the power set of speakers and $|\mathcal{N}| = 2^4$ (we assume at most four speakers). For each mapped frame-level embedding \mathbf{v}'_t , we compute its hyperbolic distance to each prototype:

$$d_{t,n} = d_{\mathbb{B}_c}(\mathbf{v}'_t, \mathbf{p}_n) \quad (5)$$

Finally, the per-speaker frame-level activity $\pi_{t,s}$ is obtained by first applying an element-wise activation function to produce a joint distribution over all classes and then marginalizing them:

$$\pi_{t,s} = \sum_{n=1}^{2^{\mathcal{N}}} b_{s,n} \sigma(-d_{t,n}), s = 1, 2, 3, 4 \quad (6)$$

where $b_{s,n} \in \{0, 1\}$ indicates whether speaker s in class n .

4.2 Speaker-Time Aware Encoder

TellWhisper adopts TS-RoPE to inject temporal and frame-level speaker activity cues into self-attention by rotating Query/Key vectors in multiple interleaved rotary subspaces.

4.2.1 Position Construction

As shown in the lower-right part of Fig. 2, for each encoder convolution layer output frame f_t , we construct a position vector consisting of one temporal index ψ_{time} and four speaker-dependent indices ψ_{spk_s} . Meanwhile, we partition the f_t 's channel dimension D into groups of 16 dimensions. Within each group, the 8 rotary pairs are assigned ψ in an interleaved manner: $[\psi_{time}, \psi_{spk_1}, \psi_{time}, \psi_{spk_2}, \psi_{time}, \psi_{spk_3}, \psi_{time}, \psi_{spk_4}]$. For the temporal position, we use the temporal index:

$$\psi_{time}(f_t) = t, \quad t \in \{0, 1, \dots, T-1\} \quad (7)$$

For the speaker-dependent indices, to capture both *within-speaker continuity* and *speaker-turn*, we define a cumulative speaker-turn counts \mathcal{C} . It first obtain a binary activity indicator with a small threshold τ (e.g., if $\pi_{t-1,s} = 0.03$ and $\pi_{t,s} = 0.8$, then $a_{t-1,s} = 0$ and $a_{t,s} = 1$):

$$a_{t,s} = \mathbb{I}[\pi_{t,s} \geq \tau], \tau = 0.1 \quad (8)$$

¹*silence*; single-speaker sets $\{1\}, \{2\}, \{3\}, \{4\}$; two-speaker overlaps $\{1, 2\}, \dots, \{3, 4\}$; three-speaker overlaps $\{1, 2, 3\}, \dots, \{2, 3, 4\}$; and $\{1, 2, 3, 4\}$.

It then detect rising edges (i.e., a speaker starts speaking means a new turn segment / turn) and accumulate them:

$$r_{t,s} = a_{t,s}(1 - a_{t-1,s}), \quad a_{0,s} = 0$$

$$C_{t,s} = \sum_{i=0}^t r_{i,s} \quad (9)$$

Finally, the speaker position index is composed of the cumulative speaker-turn counts $C_{t,s}$ and a within-turn activity:

$$\psi_{spk_s}(f_t) = C_{t,s} + \pi_{t,s} \quad (10)$$

In addition, to encourage subsequent self-attention to focus more on the *active-speaker* components in the Query, we introduce an extra, dynamic phase bias on the Query in speaker sub-spaces:

$$\psi'_{spk_s}(f_t) = \psi_{spk_s}(f_t) + (1 - \pi_{t,s}) \quad (11)$$

note we apply the bias only to Query while keeping Key unchanged.

4.2.2 TS-RoPE-Based Self-Attention

Let $\mathbf{q}_{f'_t}, \mathbf{k}_{f_t} \in \mathbb{R}^D$ denote the Query and Key vectors at frame f'_t and f_t . For the i -th rotary pair, if the pair in time region, the rotation angle is defined as:

$$\theta_{f_t,i} = \psi_{time}(f_t) \omega_i, \quad \theta_{f'_t,i} = \psi_{time}(f'_t) \omega_i \quad (12)$$

if the pair in speaker region:

$$\theta_{f_t,i} = \psi_{spk_s}(f_t) \omega_i, \quad \theta_{f'_t,i} = \psi'_{spk_s}(f'_t) \omega_i \quad (13)$$

where ω_i is the corresponding inverse frequency (all 8 rotary pairs within the same group share the same ω):

$$\omega_i = \frac{1}{10000^{\frac{2i}{D}}}, \quad i = 0, 1, \dots, \frac{D}{16} - 1 \quad (14)$$

The rotary transformation \mathcal{R} is applied simultaneously to the Query and Key:

$$\mathcal{R}(\mathbf{x}_{f_t})_i = \begin{bmatrix} x_{f_t,2i} \cos \theta_{f_t,i} - x_{f_t,2i+1} \sin \theta_{f_t,i} \\ x_{f_t,2i} \sin \theta_{f_t,i} + x_{f_t,2i+1} \cos \theta_{f_t,i} \end{bmatrix},$$

$$\mathbf{x}_{f_t} \in \{\mathbf{q}_{f'_t}, \mathbf{k}_{f_t}\} \quad (15)$$

After applying TS-RoPE, the attention weight between frames f'_t and f_t can be written as

$$Attn(f'_t, f_t) \propto \langle \mathcal{R}(\mathbf{q}_{f'_t}), \mathcal{R}(\mathbf{k}_{f_t}) \rangle \quad (16)$$

By coupling temporal positions with cumulative speaker phases, the resulting attention jointly captures temporal and speaker dynamics, yielding a fused representation E that aligns “*who*” and “*when*” cues for the subsequent Structured Content Predictor.

4.3 Structured Content Predictor

As shown in the upper-left part of Fig. 2, For the output content of TellWhisper, we adopt a segment-level structured modeling strategy. Specifically, temporally contiguous speech regions produced by the same speaker are treated as individual speech segments, each represented by an ordered sequence of tokens: $\langle spk_s \rangle$, $\langle t_{start} \rangle$, $\langle text \rangle$, and $\langle t_{end} \rangle$. All speech segments from different speakers are concatenated in chronological order to form the final target sequence. For modeling, we employ a language-model-based autoregressive framework, treating the structured representation as a unified token sequence and training it using next-token prediction. During decoding, the model generates tokens sequentially conditioned on the encoded audio representations until the end-of-sequence token $\langle EOS \rangle$ is produced.

5 Experiments

To validate the effectiveness of the proposed TellWhisper in MASR task, we conduct comprehensive experiments. In addition, to assess the reliability of speaker activity produced by Hyper-SD, we carry out comprehensive evaluations on the SD task. In this section, we describe the experimental setup from the perspectives of Datasets, Metrics, Baseline Models and Training Strategy.

5.1 Datasets

For the MASR task, we select four English multi-speaker datasets for training and evaluation. *AMI* (SDM) (Kraaij et al., 2005) and *NotSoFar* (Vinnikov et al., 2024) are collected from real-world multi-party meetings and recorded in far-field conditions, whereas *Libri2Mix* (Cosentino et al., 2020) and *LibriCSS* (Chen et al., 2020) are simulated. We also use single-utterance *LibriSpeech* (Panayotov et al., 2015) for preliminary fine-tuning before MASR training. **For the SD task**, we use six datasets for training and evaluation: *AISHHELL4* (Fu et al., 2021), *AliMeeting* (Yu et al., 2022), *AMI*, *MS-DWild* (Liu et al., 2022), *RAMC* (Yang et al., 2022), and *VoxConverse* (Chung et al., 2020), all consisting of real-world multi-speaker conversations. For

detailed statistics (speech duration, overlap duration, and number of speakers), please refer to the Appendix A.1.

5.2 Metrics

To evaluate MASR in multi-speaker settings, conventional word error rate (*WER*) is inadequate, as it fails to address speaker-permutation ambiguity and temporal misalignment. Using the Meeteval toolkit², we report four metrics: (1) *Concatenated minimum-permutation WER (CP-WER)*, measuring content accuracy with speaker attribution. (2) *Time-constrained minimum-permutation WER (TCP-WER)*, adding temporal constraints to assess consistency of content, speaker, and time. (3) *Optimal reference combination WER (ORC-WER)*, a speaker-independent WER. (4) *Time-constrained ORC-WER (TCORC-WER)*, adding temporal constraints to *ORC-WER*. For *TCP-WER* and *TCORC-WER*, we set the collar to 0.5, i.e., a small forgiveness window around reference boundaries where timing deviations are ignored.

For *SD*, we use diarization error rate (*DER*) with collar settings of 0.0 and 0.5.

5.3 Baseline Models

To comprehensively evaluate TellWhisper on the MASR task, we benchmark it against three categories of state-of-the-art baselines: (1) *Alignment-based models*, including *Pyannote3*³ + *Whisper* and *Hyper-SD* + *Whisper*, which align and integrate the outputs of speaker diarization and a single-speaker ASR model via timestamps. (2) *Separation-based model*, *Tiger* (Xu et al., 2024) + *Whisper*, which first extracts the target speaker’s speech using the high-performing speech separation model and then performs single-speaker recognition. (3) *Single-stage prediction-based model*, including *Whisper-D* (fine-tuned directly from a single-utterance ASR model), *SortFormer* (Park et al., 2024) (adding speaker posteriors to the speech-encoder outputs), *Dicow* (Polok et al., 2025) (applying speaker masks before speech encoding) and *TellWhisper-Diarizen* (replace *Hyper-SD* with *Diarizen*). For a fair comparison, all baselines are trained and fine-tuned on the same backbone as *TellWhisper*, i.e., *Whisper large-v3-turbo*⁴.

To assess the reliability of *Hyper-SD* on the speaker diarization task, we compare it with two

leading open-source models, *Pyannote3*⁵ and *Diarizen* (Han et al., 2025), both of which operate in Euclidean space. The former uses convolutional and linear layers, whereas the latter uses *WavLM*, *Conformer*, and a linear layer.

5.4 Training Strategy

We initialize *TellWhisper* with the pretrained *Whisper large-v3-turbo*⁴ and freeze the first two convolutional layers of the encoder. To match *Dicow*’s training setup, we adopt a two-stage fine-tuning strategy: we first pre-fine-tune on single-speaker speech to learn structured content prediction for a single speaker, and then fine-tune on multi-speaker conversational speech to learn structured content prediction for multiple speakers. We apply the same training pipeline to *Whisper-D* and *SortFormer*. The models are trained with token-level cross-entropy using the *AdamW* optimizer (Loshchilov and Hutter, 2017).

For *Hyper-SD*, we initialize the *WavLM* backbone with *WavLM-Large*⁶ and train on conversational data using *NLLLoss*. We optimize the hyperbolic classifier with *RiemannianAdam* (Yun and Yang, 2023) and the remaining components with *AdamW*, employing a smaller learning rate for *WavLM* and a larger one for the other modules.

6 Results and Discussions

Models	DER (↓)					
	ζ=0s		ζ=0.25s		ζ=0s	
	AMI		AISHELL4		AliMeeting	
Pyannote3 [▲]	22.60	15.41	11.96	6.27	24.40	15.67
Diarizen [▲]	13.99	9.00	9.94	4.78	13.03	5.98
Hyper-SD	13.62	8.82	9.52	4.44	10.76	4.59
Models	MSDWild		RAMC		VoxConverse	
Pyannote3 [▲]	21.73	12.25	20.91	12.97	11.18	6.81
Diarizen [▲]	12.33	5.09	11.20	6.54	9.19	5.74
Hyper-SD	12.28	4.79	10.94	6.48	8.75	5.21

Table 1: Speaker diarization results of *Hyper-SD* on conversational speech. The symbol ▲ denotes models operating in Euclidean space. ζ is the collar.

In this section, we comprehensively evaluate *TellWhisper*. We first validate the diarization capability of *Hyper-SD* and the reliability of its speaker-activity estimates. We then evaluate *TellWhisper* on MASR for jointly predicting speakers, timestamps, and transcribed content. To quantify the contribution of each TS-RoPE component, we conduct ablation studies. Finally, we visualize the

²<https://github.com/fgnt/meeteval>

³<https://github.com/yinruiqing/pyannote-whisper>

⁴<https://huggingface.co/openai/whisper-large-v3-turbo>

⁵<https://huggingface.co/pyannote/speaker-diarization-3.1>

⁶<https://huggingface.co/microsoft/wavlm-large>

Models	CP-WER (↓)				TCP-WER (↓)			
	Libri2Mix	AMI	NotSoFar	LibriCSS	Libri2Mix	AMI	NotSoFar	LibriCSS
<i>Processing: speaker diarization + single-speaker speech recognition (results alignment)</i>								
Pyannote3+Whisper [¶]	62.05	59.58	69.85	44.34	62.08	61.21	70.89	44.74
Hyper-SD+Whisper [¶]	61.23	58.51	67.22	42.51	61.25	59.62	67.84	42.68
<i>Processing: speech decoupling → single-speaker speech recognition</i>								
Tiger+Whisper [¶]	37.96	-	-	-	37.97	-	-	-
<i>Processing: multi-speaker speech recognition</i>								
Whisper-D [¶]	14.48	35.23	38.04	12.41	14.57	36.86	38.15	12.58
SortFormer [¶]	14.62	34.24	36.54	12.16	14.76	35.96	36.73	12.88
Dicow [¶]	14.34	33.57	35.22	10.62	14.35	34.02	35.64	11.33
TellWhisper-Diarizen	14.45	33.12	34.81	9.93	14.87	33.72	34.86	11.15
TellWhisper (ours)	<u>14.39</u>	32.53	34.48	9.88	<u>14.61</u>	33.47	34.51	11.06

Table 2: Multi-speaker ASR results of TellWhisper on conversational speech. CP-WER measures content + speaker, TCP-WER measures time + content + speaker. The symbol ¶ denotes absolute positional encoding.

Models	OCR-WER (↓)			
	Libri2Mix	AMI	NotSoFar	LibriCSS
Whisper-D [¶]	14.39	34.16	35.67	11.96
SortFormer [¶]	14.51	33.11	34.52	11.73
Dicow [¶]	13.34	32.83	32.20	9.43
TellWhisper-Diarizen	13.46	31.35	32.52	9.16
TellWhisper (ours)	13.32	30.72	32.31	9.14
Models	TCORC-WER (↓)			
	Libri2Mix	AMI	NotSoFar	LibriCSS
Whisper-D [¶]	14.40	35.81	34.24	12.25
SortFormer [¶]	14.55	34.57	35.21	12.42
Dicow [¶]	13.36	33.53	32.43	11.05
TellWhisper-Diarizen	13.83	32.11	32.45	10.47
TellWhisper (ours)	<u>13.67</u>	31.87	32.36	10.42

Table 3: Multi-speaker ASR results of TellWhisper on conversational speech. CP-WER measures content, TCP-WER measures time + content. The symbol ¶ denotes absolute positional encoding.

distribution of Hyper-SD class prototypes in hyperbolic space. In Appendix B, we further provide qualitative case studies on the impact of Hyper-SD’s curvature hyperparameter c on classification performance, as well as TellWhisper’s recognition performance under different overlap ratios.

6.1 Verifying the Reliability of Hyper-SD

In this experiment, we compare against Pyannote3 and Diarizen. Table 1 reports DER under two collar settings (0 s and 0.25 s). Overall, Hyper-SD attains the best DER on all datasets for both collars, indicating robust and consistent gains. In particular, both Diarizen and Hyper-SD markedly outperform Pyannote3, indicating that WavLM-based encoders can extract richer speaker-related acoustic information from speech frames than CNN-based structure. Compared with Diarizen, Hyper-SD yields the largest improvement on AliMeeting (the improvement is 2.27 when $c = 0$ s and 1.59 when $c = 0.25$ s), indicating more robust speaker separability and activity estimation in challenging real

meeting conditions. Consistent improvements are also observed on other datasets, e.g., AMI (13.99 → 13.62; 9.00 → 8.82) and AISHELL4 (9.94 → 9.52; 4.78 → 4.44). These results indicate that classifying learned speech representations in hyperbolic space is more effective than performing linear classification directly in Euclidean space. This further supports the reliability of its speaker-activity estimation, providing a more stable prior for subsequent “who speaks when” modeling in MASR.

6.2 Evaluating the Performance of Multi-Speaker Speech Recognition

In the MASR experiments, we evaluate on four datasets, and the results in Table 2 exhibit a clear hierarchy across paradigms. The “diarization + single-speaker ASR” pipeline performs worst, indicating strong sensitivity to upstream separation/alignment errors and error propagation. Tiger+Whisper reduces Libri2Mix WER to 37.96/37.97, yet still falls behind direct multi-speaker recognition. Among single-stage systems, TellWhisper achieves the best performance and TellWhisper-Diarizen the second-best on AMI, NotSoFar, and LibriCSS, consistently surpassing Dicow while also reducing TCP-WER, suggesting improved speaker attribution without compromising timestamp accuracy. TellWhisper further outperforms TellWhisper-Diarizen on all datasets (e.g., WER $-0.59/ -0.25$ on AMI), confirming the benefit of Hyper-SD. On fully overlapped Libri2Mix, our approach matches the strongest baseline, with larger gains on real meetings. This is likely due to Libri2Mix’s construction: overlap starts at time zero and each speaker has a single utterance, resulting in no speaker-turn transitions. As TS-RoPE targets speaker-aware temporal dynamics, such struc-

Models	CP-WER (\downarrow)				TCP-WER (\downarrow)			
	Libri2Mix	AMI	NotSoFar	LibriCSS	Libri2Mix	AMI	NotSoFar	LibriCSS
TellWhisper (A)	14.39	32.53	34.48	9.88	14.61	33.47	34.51	11.06
(A) w/o M_{query} (B)	15.13	35.02	36.27	10.82	15.38	35.26	37.13	12.61
(B) w/o $M_{speaker-turn}$ (C)	15.53	36.22	38.13	11.68	15.60	36.68	39.23	12.84
(C) w/o $M_{activity}$	15.48	36.84	39.54	12.32	15.50	36.89	39.63	12.75

Table 4: Ablation results of TellWhisper, where M_{query} denotes the extra angular rotation applied to the Query speaker region, $M_{speaker-turn}$ denotes cumulative speaker-turn counts, and $M_{activity}$ denotes speaker activity.

ture offers limited headroom for further WER reductions, while remaining competitive under extreme overlap.

Table 3 further corroborates this conclusion from a content-centric perspective: TellWhisper reduces OCR-WER to 30.72 / 9.14 on AMI / LibriCSS and achieves the lowest TCOCR-WER on AMI / NotSoFar / LibriCSS (31.87 / 32.36 / 10.42), with only a slight degradation relative to Dicow on Libri2Mix. Overall, TellWhisper’s advantages are most evident in real meeting and conversational scenarios with more frequent overlap and more complex speaker turns, demonstrating stronger speaker modeling and more robust temporal alignment.

6.3 Ablation Results

We ablate the design of speaker-region positional indices in TS-RoPE. As shown in Table 4, with all components enabled, TellWhisper achieves optimal performance on both CP-WER and TCP-WER. Removing the extra Query-side phase bias (w/o M_{query}) consistently degrades performance (CP-WER +0.74~2.49; TCP-WER +0.77~2.62), suggesting this Query-only phase encourages attention to emphasize active speakers, improving speaker assignment and temporal alignment. Further removing the cumulative speaker-turn counts (w/o $M_{speaker-turn}$) causes larger drops (CP-WER +1.14~3.69; TCP-WER +0.99~4.72), especially on AMI/NotSoFar, highlighting the importance of cumulative turn information for continuity and turn boundaries. When removing posterior-based activity cues in the speaker region (w/o $M_{posterior}$), performance drops most severely (NotSoFar CP-WER / TCP-WER +5.06 / +5.12), indicating posteriors are the key signal for identifying active speakers and maintaining stable alignment.

6.4 Visualization Results

As SD requires frame-level assignment to speaker classes, it primarily relies on fine-grained discriminative structure rather than an abstract-to-specific hierarchy. We therefore visualize the learned pro-

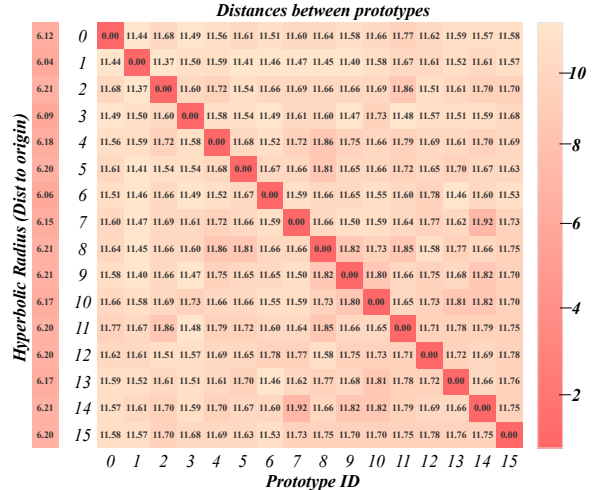


Figure 3: Visualization of hyperbolic distances among the 16 class prototypes and their distances to the origin in hyperbolic-space-based speaker activity estimation.

totypes by plotting their pairwise hyperbolic distance matrix together with each prototype’s radial distance to the origin. As shown in Fig. 3, the inter-prototype distances are largely uniform (around 11-12, right) and the radii vary within a narrow range (around 6.0-6.2, left), indicating that the prototypes are well separated and exhibit no clear hierarchical stratification.

7 Conclusion

We present TellWhisper, a unified framework for multi-speaker automatic speech recognition that couples temporal structure with speaker dynamics in the speech encoder. The core of TellWhisper is TS-RoPE, a time-speaker-aware rotary encoding that partitions Query/Key channels into temporal and speaker subspaces and applies region-specific rotations to align “when” and “who” cues in self-attention. TS-RoPE uses frame-level speaker activity to build speaker coordinates that capture within-speaker continuity and turn transitions. For reliable activity estimates, Hyper-SD performs prototype-based speaker-combination classification in hyperbolic space and derives activity from feature-

prototype distances. Experiments show TellWhisper improves recognition accuracy, speaker attribution, and time consistency, while Hyper-SD delivers robust diarization and stable activity priors. These results indicate time-speaker-aware positional modeling and geometry-aware classification effectively support multi-speaker speech understanding.

Limitations

(1) The number of supported speakers is limited. In this work, TS-RoPE is designed for 1-4 speaker scenarios, consistent with the settings in SortFormer and DiCoW. We further extend the model to accommodate more speakers in future work. **(2) The geometric consistency of Hyper-SD remains to be improved.** At present, we perform hyperbolic classification only after feature extraction, leaving the encoder and the hyperbolic classifier in mismatched embedding spaces. In future work, we investigate end-to-end hyperbolic learning that encodes features directly in hyperbolic space using hyperbolic neural networks.

Ethics Statement

(1) About the datasets: All datasets used in this work are sourced from publicly available corpora and are accessed only after obtaining approval through the official application processes on their respective websites. We use only the audio recordings and the transcripts, and the data do not contain personally identifying information or offensive content. **(2) About AI assistants:** During model development, we use AI tools to assist with code writing, and all AI-generated code is subsequently reviewed and verified by humans. During paper writing, we use AI only for grammar checking.

Acknowledgment

This research of Rui Liu was funded by the General Program (No.62476146) of the National Natural Science Foundation of China, the Young Elite Scientists Sponsorship Program by CAST (2024QNRC001), the Outstanding Youth Project of Inner Mongolia Natural Science Foundation (2025JQ011), the Key R&D and Achievement Transformation Program of Inner Mongolia Autonomous Region (2025YFHH0014), the Central Government Fund for Promoting Local Scientific and Technological Development (2025ZY0143).

References

- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, and others. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Hervé Bredin, Ruiqing Yin, Juan Manuel Coria, and others. 2020. Pyannote. audio: neural building blocks for speaker diarization. In *ICASSP*, pages 7124–7128. IEEE.
- Nan Cao, Yu-Ru Lin, Xiaohua Sun, David Lazer, and others. 2012. Whisper: Tracing the spatiotemporal process of information diffusion in real time. *IEEE transactions on visualization and computer graphics*.
- Moran Chen and Mingjiang Wang. 2024. An investigation of rotary position embedding for speech enhancement. In *Proceedings of the 2024 4th International Conference on Signal Processing and Communication Technology*, pages 44–48.
- Sanyuan Chen, Chengyi Wang, and others. 2022. Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6):1505–1518.
- Zhuo Chen, Takuya Yoshioka, Liang Lu, and others. 2020. Continuous speech separation: Dataset and analysis. In *ICASSP*, pages 7284–7288. IEEE.
- Joon Son Chung, Jaesung Huh, Arsha Nagrani, Triantafyllos Afouras, and Andrew Zisserman. 2020. Spot the conversation: speaker diarisation in the wild. *arXiv preprint arXiv:2007.01216*.
- Joris Cosentino, Manuel Pariente, Samuele Cornell, and others. 2020. Librimix: An open-source dataset for generalizable speech separation. *arXiv preprint arXiv:2005.11262*.
- Yihui Fu, Luyao Cheng, Shubo Lv, Yukai Jv, and others. 2021. Aishell-4: An open source dataset for speech enhancement, separation, recognition and speaker diarization in conference scenario. *arXiv preprint arXiv:2104.03603*.
- Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. 2018. Hyperbolic neural networks. *Advances in neural information processing systems*, 31.
- Pengcheng Guo, Xuankai Chang, Hang Lv, Shinji Watanabe, and Lei Xie. 2024. Sq-whisper: Speaker-querying based whisper model for target-speaker asr. *IEEE/ACM TASLP*.
- Jiangyu Han, Federico Landini, Johan Rohdin, Anna Silnova, and others. 2025. Leveraging self-supervised learning for speaker diarization. In *ICASSP*, pages 1–5. IEEE.
- Xiaoshuai Hao, Lei Zhou, Zhijian Huang, Zhiwen Hou, Yingbo Tang, Lingfeng Zhang, and 1 others. 2025. Mimo-embodied: X-embodied foundation model technical report. *arXiv preprint arXiv:2511.16518*.

- Xiluo He, Alexander Polok, Jesús Villalba, Thomas Thebaud, and Matthew Maciejewski. 2025. Scaling multi-talker asr with speaker-agnostic activity streams. *arXiv preprint arXiv:2510.03630*.
- Zeqian Ju, Dongchao Yang, Jianwei Yu, Kai Shen, Yichong Leng, Zhengtao Wang, Xu Tan, Zhou, and others. 2025. Mooncast: High-quality zero-shot podcast generation. *arXiv preprint arXiv:2503.14345*.
- Jiawen Kang, Lingwei Meng, Mingyu Cui, Yuejiao Wang, and others. 2025. Disentangling speakers in multi-talker speech recognition with speaker-aware ctc. In *ICASSP*, pages 1–5. IEEE.
- Wessel Kraaij, Thomas Hain, Mike Lincoln, and Wilfried Post. 2005. The ami meeting corpus. In *Proc. International Conference on Methods and Techniques in Behavioral Research*, pages 1–4.
- Tao Liu, Shuai Fan, Xu Xiang, Hongbo Song, and others. 2022. Msdwild: Multi-modal speaker diarization dataset in the wild. In *INTERSPEECH*, pages 1476–1480.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Zeyu Lu, Zidong Wang, Di Huang, Chengyue Wu, and others. 2024. Fit: Flexible vision transformer for diffusion model. *arXiv preprint arXiv:2402.12376*.
- Hao Ma, Zhiyuan Peng, Mingjie Shao, Jing Li, and Ju Liu. 2024. Extending whisper with prompt tuning to target-speaker asr. In *ICASSP*, pages 12516–12520. IEEE.
- Ivan Medennikov, Taejin Park, Weiqing Wang, He Huang, Dhawan, and others. 2025. Streaming sortformer: Speaker cache-based online speaker diarization with arrival-time ordering. *arXiv preprint arXiv:2507.18446*.
- Avik Pal, Max van Spengler, di Melendugno, and others. 2024. Compositional entailment learning for hyperbolic vision-language models. *arXiv preprint arXiv:2410.06912*.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. Librispeech: an asr corpus based on public domain audio books. In *ICASSP*, pages 5206–5210. IEEE.
- Taejin Park, Ivan Medennikov, Kunal Dhawan, and others. 2024. Sortformer: Seamless integration of speaker diarization and asr by bridging timestamps and tokens. *arXiv preprint arXiv:2409.06656*.
- Darius Petermann and Minje Kim. 2024. Hyperbolic distance-based speech separation. In *ICASSP*, pages 1191–1195. IEEE.
- Alexander Polok, Dominik Klement, Martin Kocour, Jiangyu Han, Federico Landini, and others. 2025. Dicow: Diarization-conditioned whisper for target speaker automatic speech recognition. *Computer Speech & Language*, page 101841.
- Hyorim Shin, Hanna Chung, Chaieun Park, and Soojin Jun. 2025. Enhancing the multi-user experience in fully autonomous vehicles through explainable ai voice agents. *International Journal of Human-Computer Interaction*, 41(11):6672–6686.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, and others. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Abraham A Ungar. 2001. Hyperbolic trigonometry and its application in the poincaré ball model of hyperbolic geometry. *Computers & Mathematics with Applications*, 41(1-2):135–147.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Alon Vinnikov, Amir Ivry, and others. 2024. Notssofar-1 challenge: New datasets, baseline, and tasks for distant meeting transcription. *arXiv preprint arXiv:2401.08887*.
- Zheng Weihua, Xin Huang, Zhengyuan Liu, Tarun Kumar Vangani, Bowei Zou, Xiyan Tao, Yuhao Wu, AiTi Aw, Nancy F Chen, and Roy Ka-Wei Lee. 2026. Adamcot: Rethinking cross-lingual factual reasoning through adaptive multilingual chain-of-thought. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 33863–33871.
- Hanke Xie, Haopeng Lin, and others. 2025. Soulx-podcast: Towards realistic long-form podcasts with dialectal and paralinguistic diversity. *arXiv preprint arXiv:2510.23541*.
- Kai-Tuo Xu, Feng-Long Xie, Xu Tang, and Yao Hu. 2025. Fireredasr: Open-source industrial-grade mandarin speech recognition models from encoder-decoder to llm integration. *arXiv preprint arXiv:2501.14350*.
- Mohan Xu, Kai Li, Guo Chen, and Xiaolin Hu. 2024. Tiger: Time-frequency interleaved gain extraction and reconstruction for efficient speech separation. *arXiv preprint arXiv:2410.01469*.
- Shu-Lin Xu, Yifan Sun, Faen Zhang, Anqi Xu, and others. 2023. Hyperbolic space with hierarchical margin boosts fine-grained learning from coarse labels. *Advances in Neural Information Processing Systems*, 36:71263–71274.
- Hiroyoshi Yamasaki, Jérôme Louradour, Julie Hunter, and Laurent Prévot. 2023. Transcribing and aligning conversational speech: A hybrid pipeline applied to french conversations. In *IEEE Automatic Speech Recognition and Understanding Workshop*, pages 1–6. IEEE.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, and others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

Zehui Yang, Yifan Chen, Lei Luo, Runyan Yang, Lingxuan Ye, and others. 2022. Open source magicdata-ramc: A rich annotated mandarin conversational (ramc) speech dataset. *arXiv preprint arXiv:2203.16844*.

Zengwei Yao, Liyong Guo, Xiaoyu Yang, Wei Kang, and others. 2023. Zipformer: A faster and better encoder for automatic speech recognition. *arXiv preprint arXiv:2310.11230*.

Han Yin, Yafeng Chen, Chong Deng, Luyao Cheng, and others. 2025. Speakerlm: End-to-end versatile speaker diarization and recognition with multimodal large language models. *arXiv preprint arXiv:2508.06372*.

Fan Yu, Shiliang Zhang, Yihui Fu, Lei Xie, Siqi Zheng, Zhihao Du, and others. 2022. M2met: The icassp 2022 multi-channel multi-party meeting transcription challenge. In *ICASSP*, pages 6167–6171. IEEE.

Jihun Yun and Eunho Yang. 2023. Riemannian sam: Sharpness-aware minimization on riemannian manifolds. *Advances in Neural Information Processing Systems*, 36:65784–65800.

Shucong Zhang, Titouan Parcollet, Rogier van Dalen, and Sourav Bhattacharya. 2025. Benchmarking rotary position embeddings for automatic speech recognition. *arXiv preprint arXiv:2501.06051*.

Shengkui Zhao and Bin Ma. 2023. Mossformer: Pushing the performance limit of monaural speech separation using gated single-head transformer with convolution-augmented joint self-attentions. In *ICASSP*, pages 1–5. IEEE.

Technical Appendix

In this technical appendix, we provide additional details of TellWhisper for reference, including experimental settings and supplementary results.

A More Details of Experiments

In this section, we provide additional experimental details, including the datasets, experimental setup.

A.1 Datasets

Statistics of the four MASR datasets are summarized in Table 5, including the duration breakdown of the training, validation, and test splits, as well as the proportion of overlapping speech in each dataset. Among them, Libri2Mix exhibits the highest overlap ratio, mainly because each utterance is constructed by mixing two single-sentence recordings from different speakers, resulting in overlap

Datasets	Split	Speech Duration	Overlap Duration	Max Speaker
AMI	train	65.81	8.59	4
	dev	7.69	1.06	4
	test	7.39	1.04	4
NotSoFar	train	31.15	6.80	4
	dev	13.99	3.51	4
	test	15.99	3.95	4
Libri2Mix	train	346.88	264.82	2
	dev	7.23	4.21	2
	test	2.16	1.42	2
LibriCSS	dev	1.00	0.07	4
	test	8.66	0.60	4

Table 5: Statistics of the MASR datasets, including speech duration (h), overlapped-speech duration (h), and the maximum number of speakers.

Datasets	Split	Speaker proportion			
		1	2	3	4
AMI	train	12.67	24.75	33.51	29.07
	dev	12.41	21.75	30.85	34.99
	test	14.59	23.09	32.36	29.96
NotSoFar	train	1.95	6.56	17.97	73.52
	dev	2.19	8.11	16.43	73.25
	test	3.83	8.35	24.51	63.31
Libri2Mix	train	0.00	100.00	0.00	0.00
	dev	0.00	100.00	0.00	0.00
	test	0.00	100.00	0.00	0.00
LibriCSS	dev	10.64	29.13	30.48	29.75
	test	11.85	28.64	30.68	28.83

Table 6: Speaker-count distribution of the multi-speaker ASR datasets, reporting the proportion (%) of utterances with each number of speakers in each dataset.

starting from time 0:00. In addition, to match the TS-RoPE setting in our model, we segment all datasets such that each utterance contains at most four speakers (i.e., 1-4 speakers). As shown in Table 6, each dataset includes multi-speaker utterances with different speaker-count distributions.

Statistics of the six SD datasets are reported in Table 7, including total speech duration, overlapping speech duration, and the maximum number of speakers. During Hyper-SD training, we store time-stamped supervision in *RTTM* format. Each training chunk contains 799 frames, and we additionally impose an upper bound on the number of speakers per segment (i.e., 4 speakers).

A.2 Experimental Setup

As shown in Table 8, we report the key hyperparameters of the main modules in TellWhisper. During training, we adopt different optimizers and learning rates for different components. (1) **Speaker Activ-**

Datasets	Split	Speech duration	Overlap duration	Max speaker
AISHELL4	train	97.22	87.44	7
	dev	9.36	0.76	7
	test	11.51	0.57	7
AMI	train	64.98	8.72	5
	dev	7.00	0.99	4
	test	7.29	1.06	4
AliMeeting	train	103.44	29.71	4
	dev	3.88	0.84	4
	test	9.91	2.02	4
MSDWild	train	58.67	6.84	10
	dev	6.15	0.72	7
	test	7.07	0.76	9
RAMC	train	128.68	1.20	10
	dev	8.23	0.04	2
	test	17.19	0.14	2
VoxConverse	train	16.98	0.63	20
	dev	1.93	0.08	15
	test	38.99	1.19	21

Table 7: Statistics of the speaker diarization datasets, including speech duration, overlapped-speech duration, and the maximum number of speakers.

ity Estimation (Hyper-SD). Optimizing parameters in hyperbolic space is a manifold-constrained problem with curvature, where standard Adam/AdamW (which performs Euclidean gradient updates) may lead to incorrect update directions, drifting off the manifold, and numerical instability. Therefore, for the hyperbolic speaker prototypes, we use Riemannian Adam, which performs Adam-style updates on the hyperbolic manifold, resulting in more stable optimization and faster convergence. The learning rate is set to 1×10^{-3} . For the WavLM parameters, we use AdamW with a learning rate of 2×10^{-5} ; all remaining parameters are optimized with AdamW using a learning rate of 1×10^{-3} . (2) **Speaker-Time Aware Encoder and Structured Content Predictor.** We use AdamW with a learning rate of 1×10^{-5} and $\epsilon = 1 \times 10^{-8}$. All experiments are conducted on eight NVIDIA H20 GPUs.

B More Details of Results

B.1 Hyperparameter Selection

Fig.4 presents the change in DER induced by varying the hyperbolic curvature parameter c , measured against the default $c = 1.0$ as $\Delta\text{DER} = \text{DER}(c) - \text{DER}(1.0)$, and compared under collar tolerances $\zeta \in \{0, 0.25\}$ s. We observe that across six speaker diarization datasets, $c = 1.0$ consistently yields the lowest DER under both collar settings. In contrast, $c = 0.5$ and $c = 1.5$ lead

Module	Hyperparameter	Value
<i>Frame-level Speaker Activity Estimator (Hyper-SD)</i>		
WavLM	wavlm_layer_num	25
	wavlm_feat_dim	1024
Conformer	attention_in	256
	num_head	4
	use_posi	false
Hyperbolic Projection	input_dim	256
	output_dim	128
Hyperbolic classifier	hyperbolic_dim	128
	margin	0.3
	num_classes	16
<i>Speaker-Time Aware Encoder</i>		
Tokenizer	text_n_vocab	51866
	speech_sample_rate	16000
	speech_n_mels	128
Self-Attention+MLP	d_model	1280
	attention_heads	20
	speaker_activity	0-1
	T	1500
	ffn_dim	5120
	layers (N)	32
<i>Structured Content Predictor</i>		
Decoder	attention_heads	20
	ffn_dim	5120
	layers	4
	start_token_id	50258
	eos_token_id	50257

Table 8: Partial hyperparameters of the TellWhisper.

to uniform degradation on all datasets (i.e., ΔDER is positive throughout). In particular, the degradation is most pronounced on AISHELL4; MSDWild, VoxConverse, and RAMC also show large ΔDER , suggesting that Hyper-SD is sensitive to curvature-related hyperparameters. We attribute this trend to the joint influence of c on the geometric properties of the hyperbolic manifold and its numerical behavior: under the commonly used Poincaré-ball parameterization, $c > 0$ controls the magnitude of negative curvature and the distance scale (i.e., the degree of “expansion” of the space), and as $c \rightarrow 0$, the geometry gradually degenerates to Euclidean. Therefore, a smaller c makes the space closer to Euclidean geometry, weakening the hyperbolic advantage in separating nearby classes (similar speaker representations), which may reduce inter-class/prototype separability; conversely, an excessively large c increases curvature and makes distances more sensitive to position, especially near the ball boundary, thereby amplifying numerical errors and destabilizing manifold operations and optimization. Overall, $c = 1.0$ provides a better

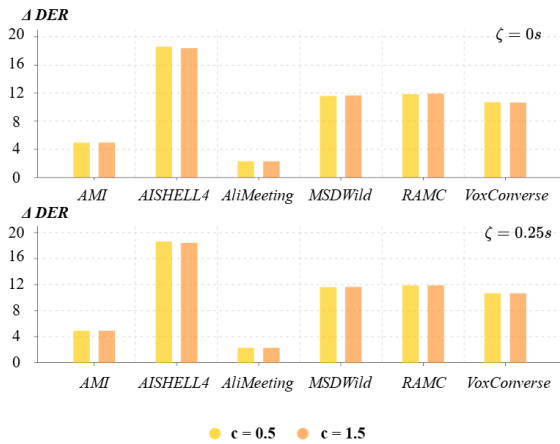


Figure 4: Comparison of DER increases relative to $c = 1.0$ under different hyperbolic-space negative-curvature parameter settings c ($\text{collar}(\zeta) = 0\text{ s} / 0.25\text{ s}$).

trade-off between representational capacity and optimization stability, and we therefore use $c = 1.0$ as the default in Hyper-SD.

B.2 Case Study

B.2.1 TellWhisper Performance on Overlapping Speech

As shown in Fig. 5 and 6, we conduct a qualitative case study on LibriCSS to examine model behav-

ior under varying overlap ratios (0%–30%), focusing on speaker assignment, temporal alignment, and content transcription. Overall, TellWhisper remains robust as overlap increases and continues to produce coherent, well-structured outputs. In terms of content, the predicted transcripts largely preserve the semantics of the ground truth, with mismatches typically limited to occasional word-level substitutions in highly overlapped regions. Regarding temporal alignment, the model generally provides reasonable start/end boundaries. Higher overlap may lead to slightly finer-grained segmentation or minor boundary shifts, yet the overall timing remains well aligned. For speaker attribution, predictions are consistently accurate under low-to-moderate overlap, while the few confusions observed at higher overlap are mostly localized around overlap windows and do not substantially disrupt the global conversational structure. Taken together, these visualizations suggest that although heavy overlap increases local ambiguity, our TellWhisper maintains strong performance across speaker, time, and content dimensions, demonstrating good robustness under challenging multi-speaker conditions.

<i>Overlap ratio: 0%</i>
<i>Ground-truth transcript:</i>
<p>Spk-0 < 00.00 > < 03.36 > husband the next thing to a wife</p> <p>Spk-1 < 06.26 > < 10.35 > can you imagine why buckingham has been so violent i suspect</p> <p>Spk-0 < 13.26 > < 15.79 > cried the ladies whose departure had been fixed</p> <p>Spk-2 < 18.70 > < 29.92 > i can set to work now to remember things i never remembered before such as what i had to eat for breakfast this morning and it can hardly be wholly habit that enables me to do this</p>
<i>Predicted transcript:</i>
<p>Spk-0 < 00.00 > < 03.32 > husband the next thing to a wife</p> <p>Spk-1 < 06.32 > < 10.26 > can you imagine why buckingham has been so violent i suspect</p> <p>Spk-0 < 13.28 > < 15.77 > cried the ladies whose departure had been fixed</p> <p>Spk-2 < 18.71 > < 22.76 > i can set to work now to remember things i never remembered before</p> <p>Spk-2 < 22.91 > < 25.84 > such as what i had to eat for breakfast this morning</p> <p>Spk-2 < 26.06 > < 29.91 > and it can hardly be holy habit that enables me to do this</p>
<i>Overlap ratio: 10%</i>
<i>Ground-truth transcript:</i>
<p>Spk-0 < 00.00 > < 01.70 > that is what you would like to be doing is it</p> <p>Spk-1 < 01.09 > < 13.81 > then they sped in great haste for the door and the goat gave a final butt that sent the row of royal ladies all diving into the corridor in another tangle whereupon they shrieked in a manner that terrified everyone within sound of their voices</p> <p>Spk-2 < 14.15 > < 22.01 > another preacher after reproaching him to his face with his misgovernment ordered this psalm to be sung</p> <p>Spk-0 < 21.24 > < 28.38 > but the windows are patched with wooden panes and the door i think is like the gate it is never opened</p>
<i>Predicted transcript:</i>
<p>Spk-0 < 00.00 > < 01.68 > that is what you would like to be doing is it</p> <p>Spk-1 < 01.00 > < 13.82 > then they sped in great haste for the door and the goat gave a final butt that sent the row of royal ladies all diving into the corridor in another tangle whereupon they shrieked in a manner that terrified every one within sound of their voices</p> <p>Spk-2 < 14.00 > < 21.94 > another preacher after reproaching him to his face with his misgovernment ordered this psalm to be sung</p> <p>Spk-0 < 21.31 > < 24.01 > but the windows are patched with wooden panes</p> <p>Spk-0 < 24.37 > < 26.72 > and the door i think is like the gate</p> <p>Spk-1 < 27.36 > < 28.35 > it is never opened</p>

Figure 5: Example transcripts on LibriCSS at 0% and 10% overlap.

<i>Overlap ratio: 20%</i>
<i>Ground-truth transcript:</i>
<p>Spk-0 < 00.00 > < 03.18 > i suppose it is the wet season will you have to cut them too</p> <p>Spk-1 < 02.20 > < 16.54 > but they dragged him out of the room and up the stairs into the loft and here in a dark corner where no daylight could enter they left him</p> <p>Spk-0 < 15.48 > < 19.27 > she asked impulsively i did not believe you could persuade her father</p> <p>Spk-2 < 18.26 > < 23.79 > since christ was given for our sins it stands to reason that they cannot be put away by our own efforts</p> <p>Spk-3 < 23.55 > < 25.54 > why should he not be as other men</p> <p>Spk-2 < 25.16 > < 28.18 > we think that by some little work or merit we can dismiss sin</p> <p>Spk-3 < 27.76 > < 29.85 > cotton she paused</p>
<i>Predicted transcript:</i>
<p>Spk-0 < 00.00 > < 03.10 > i suppose it is the wet season will you have to cut them two</p> <p>Spk-1 < 01.00 > < 16.48 > night but they dragged him out of the room and up the stairs into the loft and here in a dark corner where no daylight could enter they left him</p> <p>Spk-0 < 15.54 > < 19.24 > she asked impulsively i did not believe you could persuade her father</p> <p>Spk-2 < 18.00 > < 23.64 > since christ was given for our sins it stands to reason that they could not be put away by our own efforts</p> <p>Spk-3 < 23.60 > < 25.50 > why should he not be as other men</p> <p>Spk-2 < 25.20 > < 28.14 > we think that by some little work or merit we can dismiss some</p> <p>Spk-3 < 27.80 > < 28.58 > cotton</p> <p>Spk-3 < 29.08 > < 29.83 > she paused</p>
<i>Overlap ratio: 30%</i>
<i>Ground-truth transcript:</i>
<p>Spk-0 < 00.00 > < 07.13 > uncas who had already approached the door in readiness to lead the way now recoiled and placed himself once more in the bottom of the lodge</p> <p>Spk-1 < 05.09 > < 14.93 > the free state men clung to their prairie towns and prairie ravines with all the obstinacy and courage of true defenders of their homes and firesides</p> <p>Spk-2 < 11.65 > < 28.78 > i see a quantity of chairs for hire at the rate of one sou men reading the newspaper under the shade of the trees girls and men breakfasting either alone or in company waiters who were rapidly going up and down a narrow staircase hidden under the foliage</p>
<i>Predicted transcript:</i>
<p>Spk-0 < 00.00 > < 07.14 > unkus who had already approached the door in readiness to lead the way now recoiled and placed himself once more in the bottom of the lodge</p> <p>Spk-1 < 05.06 > < 08.82 > the three state man clung to their prairie towns and prairie ravines</p> <p>Spk-1 < 09.23 > < 14.92 > with all the obstinacy and courage of true defenders of their homes and firesides</p> <p>Spk-2 < 11.68 > < 15.34 > i see a quantity of chairs for hire at the rate of one sou</p> <p>Spk-2 < 16.16 > < 18.96 > men reading the newspaper under the shade of the trees</p> <p>Spk-2 < 19.56 > < 23.14 > girls and men breakfasting either alone or in company</p> <p>Spk-2 < 23.68 > < 28.68 > waiters who were rapidly going up and down a narrow staircase hidden under the foliage</p>

Figure 6: Example transcripts on LibriCSS at 20% and 30% overlap.