
Benchmarking Diarization Models

Luca A. Lanzendörfer¹ Florian Grötschla¹ Cesare Blaser¹ Roger Wattenhofer¹

Abstract

Speaker diarization is the task of partitioning audio into segments according to speaker identity, answering the question of “who spoke when” in multi-speaker conversation recordings. While diarization is an essential task for many downstream applications, it remains an unsolved problem. Errors in diarization propagate to downstream systems and cause wide-ranging failures. To this end, we examine exact failure modes by evaluating five state-of-the-art diarization models, across four diarization datasets spanning multiple languages and acoustic conditions. The evaluation datasets consist of 196.6 hours of multilingual audio, including English, Mandarin, German, Japanese, and Spanish. Overall, we find that PyannoteAI achieves the best performance at 11.2% DER, while DiariZen provides a competitive open-source alternative at 13.3% DER. When analyzing failure cases, we find that the primary cause of diarization errors stem from missed speech segments followed by speaker confusion, especially in high-speaker count settings.

1. Introduction

Speaker diarization, the task of finding time stamp and speaker ID information for spoken conversations, remains one of the most challenging problems in speech processing. The difficulty arises from overlapping speech segments, varying numbers of speakers, and heterogeneous recording environments (Park et al., 2022). Accurate speaker diarization has become increasingly critical for automated meeting transcriptions, voice chat analysis, broadcast content processing, and as preprocessing for automatic speech recognition (Park et al., 2022).

Contemporary diarization systems follow three primary paradigms. Modular approaches decompose the task into specialized subtasks by first detecting speech through voice

activity detection (VAD), followed by segmentation, embedding extraction, and speaker clustering (Park et al., 2022; Landini et al., 2022; Bredin et al., 2020). Various approaches additionally introduce overlap-aware segmentation (Plaquet & Bredin, 2023; Bredin & Laurent, 2021). Although this modularity enables plug-and-play replacement, error propagation from previous modules (e.g., errors in VAD) are a significant drawback of these pipelines. End-to-end neural diarization (EEND) formulates the diarization problem as a multi-label sequence classification with joint optimization (Fujita et al., 2019; Medennikov et al., 2020; Park et al., 2025), allowing explicit overlap modeling but facing scalability challenges with high speaker counts and long recordings. Recent EEND developments include streaming capabilities through mechanisms such as the Arrival-Order Speaker Cache (Medennikov et al., 2025). Hybrid approaches combine these paradigms, applying local EEND processing on audio chunks with global clustering (Kinoshita et al., 2021), balancing overlap handling with scalability.

Recent benchmarking studies (Pacheco et al., 2025; Aperdanner et al., 2024) have provided comparative insights into diarization system performance. However, comprehensive evaluation across diverse languages and recording settings using readily accessible models remains limited. To this end, we evaluate five diarization systems across four datasets spanning five languages and recording conditions (i.e., meeting recordings, phone calls, and in-the-wild audio). The tested diarization models consist of four open-source state-of-the-art models representing different paradigms, and one commercial API.

This work presents a comparative evaluation of speaker diarization models. Our evaluation addresses two questions. Firstly, the practical question of “given an audio recording in a specific language and acoustic condition, which available model achieves the best diarization performance?” and secondly, with future improvements in mind: What are the exact causes of failure for current state-of-the-art diarization models? By evaluating the performance of these models through multilingual analysis, speaker scalability, and computational efficiency, we provide users with practical guidance across different deployment scenarios, and highlight specific areas of improvement for the next generation of diarization models.

¹ETH Zurich, Switzerland. Correspondence to: Luca A. Lanzendörfer <lanzendoerfer@ethz.ch>.

2. Methodology

Our evaluation adopts an out-of-the-box methodology, testing each system as documented without parameter tuning or modifications, providing a baseline assessment of plug-and-play deployment performance. All audio was standardized to mono-channel 16 kHz sampling rate for consistent input across models. All experiments were conducted on our internal cluster using NVIDIA RTX A6000 GPUs except for pyannotAI, where we used the provided cloud-based API.

2.1. Models

Pyannote. A widely-used open-source modular diarization pipeline (Bredin, 2023; Plaquet & Bredin, 2023). The system performs diarization by chunking audio into overlapping segments, applying ResNet-based neural segmentation to detect speech regions, extracting speaker embeddings, and clustering them to assign speaker identities. We use the 3.1 version provided on Hugging Face.¹ PyannoteAI² is a cloud-based commercial service based on the Pyannote “precision-2” model, although its exact details of implementation are not publicly documented.

Sortformer (SF). An end-to-end neural diarization model (Park et al., 2025) using a transformer architecture with arrival time sorting, and optimized for up to four speakers. The authors mention that the model achieves degraded performance for higher speaker counts and that the model suffers from computational limitations for long-form audio processing, due to the quadratic memory scaling of the self-attention mechanism. We use the Hugging Face version in our experiments.³

Sortformer v2 (SF v2). A recent improvement over Sortformer, Sortformer v2 (Medennikov et al., 2025) consists of an end-to-end neural diarization model featuring Arrival-Order Speaker Cache (AOSC), designed for improved long-form audio processing capabilities. AOSC is a fixed-length buffer that stores per-speaker embeddings in arrival order while retaining the most recent and high-confidence frames, which helps resolve between-chunk speaker permutation automatically and enables robust low-latency streaming diarization.⁴

DiariZen. A hybrid diarization approach (Han et al., 2025a;b;c), combining end-to-end neural diarization processing with WavLM (Chen et al., 2022) for speaker embeddings and clustering components from Pyannote.⁵

¹ <https://huggingface.co/pyannote/speaker-diarization-3.1>

² <https://www.pyannote.ai>

³ https://huggingface.co/nvidia/diar_sortformer_4spk-v1

⁴ https://huggingface.co/nvidia/diar_streaming_sortformer_4spk-v2

⁵ <https://huggingface.co/BUT-FIT/diariZen-wavlm-large-s80-md>

Table 1. Speaker count distribution (in hours) by dataset. CALLHOME (CH) contains mostly two-speaker phone conversations. VoxConverse (VC) has 43.6 hours of audio with more than 5 speakers, with some samples reaching 15+ active speakers, coming from TV program recordings and public press events. AMI and ALI contain mostly meeting recordings with 4 active speakers.

#Speakers	CH	VC	AMI	ALI
1	0	3.2	0	0
2	74.2	7.9	0	5.9
3	16.2	4.4	0.8	2.7
4	6.0	4.7	17.9	6.4
5+	2.6	43.6	0	0

2.2. Model Adaptations & Evaluation Metric

Sortformer chunking. Sortformer (SF) requires chunking for recordings exceeding 12 minutes (specific to the hardware we used and as outlined by the authors), due to memory scaling constraints in its attention-based architecture. We implemented audio chunking by splitting recordings into 12-minute chunks, with boundaries aligned with silence gaps predicted with SileroVAD (Team, 2024), in order to preserve speech continuity and conversational dynamics. Only AMI and ALI datasets required chunking, which maintained stable speaker count and overlap characteristics with minimal deviation (0.5%).

Sortformer v2 adaptations. Sortformer v2 (SF v2) is evaluated on both the chunked audio that was processed for use by Sortformer, and the full original audio used by other models. We denote the model evaluated on full-length audio as Sortformer v2-streaming (SF v2-stream).

After processing, all model outputs were parsed into a standardized representation consisting of temporal segments with start times, end times, and speaker identifiers, allowing for consistent downstream evaluation.

Diarization error rate (DER) as evaluation metric. We evaluate all models using DER, the standard metric for speaker diarization performance (Pacheco et al., 2025), quantifying the fraction of time attributed to incorrect labels. DER decomposes into three error components: missed speech (failure to detect actual speech regions), false alarm (detection of speech in non-speech regions), and speaker confusion (incorrect speaker assignment). We implement DER using pyannote.metrics⁶ with a 0.25 second collar and skip_overlap=False to evaluate performance on overlapping speech regions where diarization systems typically struggle most.

⁶ <https://github.com/pyannote/pyannote-metrics>




Figure 1. Diarization error rate across models showing missed speech (red), false alarm (orange), and speaker confusion (blue) components. We find that Sortformer v2-streaming, DiariZen, and the commercial PyannoteAI models perform the best overall. The numbers above the bar represent the total DER (%), lower is better. We find that missed speech errors are the most common source of errors, with the exception of Sortformer, which has a higher speaker confusion error than other diarization models.

2.3. Datasets

CALLHOME Corpus (CH). A widely used multilingual telephone conversation dataset⁷ spanning five languages: English (20.3h), Mandarin (20.3h), Japanese (18.7h), German (18.4h), and Spanish (21.3h). CH contains two speakers per session, and an average of 12.58% overlapped speech.

VoxConverse v0.3 (VC). An English language conversation dataset (Chung et al., 2020) containing 63.8 hours of speech, where samples were extracted from YouTube videos containing between 2 and more than 20 speakers per recording. VoxConverse is widely used for high speaker count evaluation and in-the-wild acoustic conditions including background noise, varying microphone distances, and uncontrolled recording environments. VoxConverse contains an average 3.52% overlap speech. We evaluated the models on both the “dev” and “test” split.

AMI Meeting Corpus (AMI). A dataset (Carletta, 2007) containing English language scenario-based business meetings with exactly 4 participants per session (18.7 hours total). Used as a widely adopted benchmark and because professional meeting environments constitute a core application domain for speaker diarization. AMI contains on average 15.94% of overlap speech. We evaluated the models on both “dev” and “test” split.

AliMeeting (ALI). Dataset (Yu et al., 2022) contains 15 hours of Mandarin meeting scenarios with structure similar to AMI, featuring high speech overlap averaging 19%, with some samples exceeding 50%. We use ALI to evaluate Mandarin in addition to CALLHOME and challenging overlap conditions. We evaluated the models on both “eval” and “test” split.




Figure 2. We visualize the distribution of speech segment lengths found over all datasets (left). For each diarization model we visualize in which segment length missed speech occurred (right). We find that models exhibit mostly unbiased error patterns across segment lengths, matching dataset composition, indicating that the failures are not from missing short segments but rather from inaccurate onset and end timestamp information across all segment length types.

3. Results

3.1. General Performance

Figure 1 presents DER and its components across all model-dataset combinations, revealing that pyannoteAI achieves the best overall performance (DER = 11.2%). DiariZen produces competitive performance (DER = 13.3%) as well, with particularly strong results on VoxConverse (DER = 5.2%), while Sortformer v2-streaming demonstrates exceptional performance on ALI (DER = 7.0%), demonstrating that certain diarization models seem to be more suited for different conversational settings.

In terms of computational efficiency (cf. Table 4), we find that Sortformer’s architecture delivers exceptional speed (regardless of the version), with SF v2 achieving the fastest average processing (RTF = 214.3x) across all datasets, also

⁷ <https://huggingface.co/datasets/talkbank/callhome>

Benchmarking Diarization Models

Table 2. DER (%) aggregated by language. SF stands for Sortformer, SF v2 for Sortformer v2 with chunked audio, and SF v2-stream for Sortformer v2-streaming.

Model	Zho	Eng	Deu	Jpn	Spa
DiariZen	10.1	7.0	11.6	15.6	19.1
SF	13.1	15.5	11.1	16.5	21.9
SF v2	9.2	15.3	9.6	12.7	21.1
SF v2-stream	9.4	14.1	9.6	12.7	21.1
pyannote	19.8	11.5	19.0	28.8	27.3
pyannoteAI	10.0	6.6	8.3	13.8	14.3

Table 3. DER (%) aggregated by speaker count. For each column, we compute DER by averaging all samples with the corresponding number of speakers and averaging over them. Bold numbers indicate the best-performing system for each speaker count.

Model	1 spk	2 spk	3 spk	4 spk	5+ spk
DiariZen	2.3	11.4	10.3	12.7	7.1
SF	1.5	11.5	14.7	21.3	23.9
SF v2	4.7	10.1	14.3	16.7	22.7
SF v2-stream	4.7	10.4	14.1	13.2	22.7
pyannote	3.2	19.9	19.8	17.1	10.6
pyannoteAI	2.7	9.9	9.1	10.1	6.6

enabling real-time deployment through streaming (Medenikov et al., 2025). This represents a significant advantage for latency-critical applications, while still maintaining competitive diarization error rates.

Analyzing diarization failure modes (cf. Figure 1), we find that missed speech constitutes the dominant failure case across all diarization models, especially in meeting scenarios such as those covered in AMI and ALI. Additionally, we find that on VoxConverse all Sortformer versions exhibit increased speaker confusion errors. As VC contains on average higher average speaker counts (cf. Table 1) this reflects the documented 4-speaker limit and highlights out-of-distribution generalization gaps.

Additionally, we analyze missed speech segments to understand error characteristics. All models show consistent average missed segment durations of roughly 350ms. Analysis of ground truth segments contributing to these errors (cf. Figure 2) reveals that models exhibit minimal bias toward specific segment lengths, with the exception for very short segments (< 0.5s), which account for less than 5% of errors. The brief duration of missed speech segments predominantly combined with the distribution highlighted in Figure 2 shows that the failure is in boundary precision of longer segments, rather than completely failing to detect shorter segments. This finding suggests that focusing on the precision of speech onset and end timestamps would significantly improve diarization performance.

Table 4. Real-Time Factor (RTF, higher is better) computed as audio length divided by processing time. Processing time is only based on model inference time (without data or model loading). We do not report PyannoteAI’s RTF, as we do not have access to model inference time through the API.

Model	ALI	AMI	CH	VC	Avg
DiariZen	20.4	20.4	20.4	19.6	20.2
SF	166.7	160.7	177.9	153.4	164.7
SF v2	223.6	223.7	212.5	197.4	214.3
SF v2-stream	230.6	197.4	212.5	197.4	209.5
pyannote	45.5	47.6	43.5	43.5	45.0

3.2. Per-Language Performance

Table 2 shows cross-lingual performance patterns across five languages. PyannoteAI achieves the best performance on English (DER = 6.6%), German (DER = 8.3%), and Spanish (DER = 14.3%), while SF v2 and SF v2-stream both excel on Mandarin (DER = 9.2%) and Japanese (DER = 12.7%), suggesting different models have distinct cross-lingual strengths. English and Mandarin appear to be generally the languages where diarization performs the best, likely due to more training data availability, while Spanish presents the most challenging condition out of the evaluated languages.

3.3. Per-Speaker Count Performance

Table 3 reveals speaker count scalability patterns. PyannoteAI demonstrates the best scores in terms of scalability, maintaining relatively stable performance from 2 speakers (DER = 9.9%) to 5+ speakers (DER = 6.6%). SF v2 and SF v2-stream show improved speaker handling compared to SF, with SF v2-stream exhibiting better performance on 4-speaker scenarios (DER = 13.2% compared with SF’s DER = 21.3%), though all Sortformer-based models still experience degradation with increasing speaker counts due to their 4-speaker limit. DiariZen’s hybrid approach demonstrates particular effectiveness in high-speaker scenarios (5+: DER = 7.1%), with its strategy of processing local audio chunks with EEND from WavLM (Chen et al., 2022).

4. Conclusion

This work evaluated speaker diarization systems across four datasets containing five languages. Cross-lingual analysis indicates that annotated data scarcity for certain languages is a dominant driver in lower diarization performance, and analyzing failure cases, we found missed speech detection as the dominant failure mode across all models. Missed speech was found to be mainly caused by speech onset and end timing detection errors, highlighting possible future improvements for diarization models.

Impact Statement

This work aims to advance speaker diarization by providing practical guidance on the strengths, limitations, and failure modes of current systems across languages and recording conditions. Improved diarization can support applications such as meeting transcription, accessibility, and speech analytics. Because diarization errors can propagate to downstream systems and disproportionately affect underrepresented languages or acoustic conditions, we encourage careful validation before deployment.

References

- Aperdannier, R., Schacht, S., and Piazza, A. Systematic evaluation of online speaker diarization systems regarding their latency, 2024. URL <https://arxiv.org/abs/2407.04293>.
- Bredin, H. pyannote.audio 2.1 speaker diarization pipeline: principle, benchmark, and recipe. In *Proc. INTERSPEECH 2023*, 2023.
- Bredin, H. and Laurent, A. End-To-End Speaker Segmentation for Overlap-Aware Resegmentation. In *Interspeech 2021*, pp. 3111–3115, 2021. doi: 10.21437/Interspeech.2021-560.
- Bredin, H., Yin, R., Coria, J. M., Gelly, G., Korshunov, P., Lavechin, M., Fustes, D., Titeux, H., Bouaziz, W., and Gill, M.-P. Pyannote.Audio: Neural Building Blocks for Speaker Diarization. In *ICASSP 2020 – 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7124–7128, 2020. doi: 10.1109/ICASSP40776.2020.9052974.
- Carletta, J. Unleashing the killer corpus: experiences in creating the multi-everything ami meeting corpus. *Language Resources and Evaluation*, 41(2):181–190, 2007. ISSN 1574-020X. doi: 10.1007/s10579-007-9040-x.
- Chen, S., Wang, C., Chen, Z., Wu, Y., Liu, S., Chen, Z., Li, J., Kanda, N., Yoshioka, T., Xiao, X., Wu, J., Zhou, L., Ren, S., Qian, Y., Qian, Y., Wu, J., Zeng, M., Yu, X., and Wei, F. Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6):1505–1518, October 2022. ISSN 1941-0484. doi: 10.1109/jstsp.2022.3188113. URL <http://dx.doi.org/10.1109/JSTSP.2022.3188113>.
- Chung, J. S., Huh, J., Nagrani, A., Afouras, T., and Zisserman, A. Spot the conversation: speaker diarisation in the wild. In *Proc. Interspeech 2020*, pp. 299–303, 2020. doi: 10.21437/Interspeech.2020-1893.
- Fujita, Y., Kanda, N., Horiguchi, S., Xue, Y., Nagamatsu, K., and Watanabe, S. End-to-End Neural Speaker Diarization with Self-attention. In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pp. 296–303, 2019. doi: 10.1109/ASRU46091.2019.9003959.
- Han, J., Landini, F., Rohdin, J., Silnova, A., Diez, M., and Burget, L. Leveraging Self-Supervised Learning for Speaker Diarization. In *ICASSP 2025 – 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5, 2025a. doi: 10.1109/ICASSP49660.2025.10889475.
- Han, J., Landini, F., Rohdin, J., Silnova, A., Diez, M., Cernocky, J., and Burget, L. Fine-tune before structured pruning: Towards compact and accurate self-supervised models for speaker diarization. *arXiv preprint arXiv:2505.24111*, 2025b.
- Han, J., Pálka, P., Delcroix, M., Landini, F., Rohdin, J., Cernocký, J., and Burget, L. Efficient and generalizable speaker diarization via structured pruning of self-supervised models. *arXiv preprint arXiv:2506.18623*, 2025c.
- Kinoshita, K., Delcroix, M., and Tawara, N. Advances in Integration of End-to-End Neural and Clustering-Based Diarization for Real Conversational Speech. In *Interspeech 2021*, pp. 3565–3569, 2021. doi: 10.21437/Interspeech.2021-1004.
- Landini, F., Profant, J., Diez, M., and Burget, L. Bayesian HMM clustering of x-vector sequences (VBx) in speaker diarization: theory, implementation and analysis on standard tasks. *Computer Speech & Language*, 71:101254, 2022. doi: 10.1016/j.csl.2021.101254.
- Medennikov, I., Korenevsky, M., Prisyach, T., Khokhlov, Y., Korenevskaya, M., Sorokin, I., Timofeeva, T., Mitrofanov, A., Andrusenko, A., Podluzhny, I., Laptev, A., and Romanenko, A. Target-speaker voice activity detection: A novel approach for multi-speaker diarization in a dinner party scenario. In *Interspeech 2020*, interspeech_2020. ISCA, October 2020. doi: 10.21437/interspeech.2020-1602. URL <http://dx.doi.org/10.21437/Interspeech.2020-1602>.
- Medennikov, I., Park, T., Wang, W., Huang, H., Dhawan, K., Wang, J., Balam, J., and Ginsburg, B. Streaming sortformer: Speaker cache-based online speaker diarization with arrival-time ordering, 2025. URL <https://arxiv.org/abs/2507.18446>.
- Pacheco, E., Orhon, A., Durmus, B., Munyampirwa, B., and Leonov, A. Sdbench: A comprehensive benchmark suite for speaker diarization, 2025. URL <https://arxiv.org/abs/2507.16136>.

Park, T., Medennikov, I., Dhawan, K., Wang, W., Huang, H., Koluguri, N. R., Puvvada, K. C., Balam, J., and Ginsburg, B. Sortformer: A Novel Approach for Permutation-Resolved Speaker Supervision in Speech-to-Text Systems. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, 2025.

Park, T. J., Kanda, N., Dimitriadis, D., Han, K. J., Watanabe, S., and Narayanan, S. A review of speaker diarization: Recent advances with deep learning. *Computer Speech & Language*, 72:101317, 2022. ISSN 0885-2308. doi: <https://doi.org/10.1016/j.csl.2021.101317>. URL <https://www.sciencedirect.com/science/article/pii/S0885230821001121>.

Plaquet, A. and Bredin, H. Powerset multi-class cross entropy loss for neural speaker diarization. In *Proc. INTERSPEECH 2023*, 2023.

Team, S. Silero vad: pre-trained enterprise-grade voice activity detector (vad), number detector and language classifier. <https://github.com/snakers4/silero-vad>, 2024.

Yu, F., Zhang, S., Fu, Y., Xie, L., Zheng, S., Du, Z., Huang, W., Guo, P., Yan, Z., Ma, B., Xu, X., and Bu, H. M2MeT: The ICASSP 2022 multi-channel multi-party meeting transcription challenge. In *Proc. ICASSP. IEEE*, 2022.