

Adaptive Turn-Taking for Real-time Multi-Party Voice Agents

Soumyajit Mitra¹, Prabhat Pandey¹, Abhinav Jain¹, Shanmukha Sahith², K V Vijay Girish¹

¹ Amazon AGI

² IIT Kharagpur, India

ssomit@amazon.com, panprabh@amazon.com, jabhin@amazon.com, kvvijayg@amazon.com

Abstract

Turn-taking in multi-party spoken conversations remains a fundamental challenge for voice-based agents, particularly under dynamic floor competition and varying user expectations. We propose *ModeratorLM*, a role-playing voice agent that conditions turn-taking behavior on an explicitly assigned role in multi-party settings. The system is built on a speech large language model operating in chunk-wise streaming manner. We further introduce a reasoning-augmented variant that incorporates chain-of-thought reasoning over conversational context and the assigned role. We construct *RolePlayConv*, a large-scale synthetic dataset of spoken multi-party conversations with diverse assistant roles. Experiments on real-world meeting data and *RolePlayConv* show improved turn-taking precision by over 40% and recall by more than 70%, while substantially reducing false-positive interruptions compared to non-role-conditioned baselines.

Index Terms: Multi-party Conversation, Speech LLMs, Turn-taking, Role-playing Voice Agents

1. Introduction

Recent advances in large language models (LLMs) have driven rapid progress in the development of voice-based conversational agents [1, 2, 3, 4]. Modern spoken dialogue systems typically combine low-latency streaming speech processing modules with a core conversational component responsible for dialogue management and response generation. The emergence of full-duplex agents, capable of listening and speaking simultaneously, has further advanced voice-based conversational systems by enabling more natural interaction patterns such as backchannel responses and user barge-ins [5, 6, 7]. In dyadic (two-party) conversations, turn-taking is commonly governed by acoustic and semantic cues such as pause duration, silence detection, and sentence completion markers. Recent benchmarks have formalized the evaluation of turn-taking behavior in dyadic settings [8, 9]. However, these assumptions do not generalize well to multi-party conversations, where interactions involve overlapping speech, dynamic floor competition, and negotiated turn allocation among participants [10]. In such settings, a conversational agent must continuously decide not only *when* to speak, but also *whether* to intervene at all. At the same time, user expectations for conversational agents vary across contexts, ranging from passive listeners to active facilitators [11]. This has motivated research on role-playing language agents (RLPA), where assistants adopt explicit personas or roles [12, 13]. Prior work has primarily focused on text-based dialogue settings, modeling linguistic style [14], character traits [15], and decision-making patterns [16]. However, the effect of an assigned role on real-time turn-taking behavior in voice-

based, multi-party interaction has been largely overlooked.

In this work, we introduce *ModeratorLM*, a role-playing voice agent for multi-party conversations that conditions both turn-taking decisions and response generation on an explicitly assigned role. Our system is built on a speech LLM that processes multi-speaker audio streams and corresponding text hypotheses in chunk-wise streaming manner. To our knowledge, this is the first role-conditioned voice agent designed for multi-party conversational settings. To support training, we construct a large-scale synthetic corpus of multi-party conversations spanning diverse assistant roles using off-the-shelf LLMs and a text-to-speech (TTS) model. Experimental results show that *ModeratorLM* achieves substantially better alignment with the assigned roles than non-role-conditioned baselines, improving both role-consistent turn-taking and response generation. We further introduce *ModeratorLM-Think*, which incorporates chain-of-thought reasoning [17] over conversational context and the assigned role, yielding additional improvements in role-aligned turn-taking behavior.

2. Methodology

2.1. ModeratorLM: System Architecture

ModeratorLM consists of a speech encoder and a backbone LLM. The speech encoder processes each incoming audio chunk independently and produces chunk-level embeddings. These embeddings are projected into the LLM embedding space via a trainable linear projection layer, following prior work [18, 19]. The multi-channel audio inputs are downmixed into a single-channel signal before encoding. To support real-time interaction, chunk-level speech embeddings are appended sequentially to the LLM context in a streaming fashion, similar to [4, 20, 21]. However, unlike approaches that rely on a separate voice activity detection module to determine turn boundaries [4], our system delegates turn-taking decisions entirely to the speech LLM itself. To improve robustness to variability in chunk duration, we adopt dynamically sized training chunks. For each chunk, the corresponding text transcription with speaker annotations is provided alongside the speech embeddings as input to the LLM.

For each input audio chunk, the LLM is trained to produce one of two output types: (i) *Turn-taking + Response*, where the model emits a control token indicating that the assistant should take the floor, followed by the assistant’s textual response; or (ii) *No-turn*, represented by an empty sequence, indicating that no turn-taking action is taken for the current chunk. In addition, we train a reasoning variant, referred to as *ModeratorLM-Think*, in which the LLM performs chain-of-thought reasoning at potential turn-taking points before emitting its decision. Figure 1 illustrates an example input-output sequence for the LLM of

Model	NOTSOFAR-1						RolePlayConv					
	@P	@R	@F1	@A	@FP	@RM	@P	@R	@F1	@A	@FP	@RM
Moshi	0.14	0.10	0.11	0.21	0.66	–	0.15	0.34	0.21	0.50	0.47	–
MP-Baseline	0.58	0.33	0.38	0.69	0.05	–	0.40	0.48	0.42	0.67	0.14	–
ModeratorLM	0.77	0.51	0.57	0.77	0.01	0.08	0.71	0.57	0.61	0.76	0.05	0.14
ModeratorLM-Think	0.81	0.74	0.76	0.86	0.01	0.02	0.79	0.82	0.79	0.91	0.03	0.03

Table 2: Performance comparison across different models on NOTSOFAR-1 and RolePlayConv test sets. @P, @R, @F1, @A, @FP, and @RM denote precision, recall, F1-score, macro-accuracy, false-positive rate, and reactive miss rate, respectively. Reactive miss rate (@RM) is not applicable for Moshi and MP-Baseline, as these models do not support role configuration.

ensure that a fraction of chunks terminate at speaker boundaries, as these correspond to natural turn-taking points.

Our training pipeline consists of three stages:

Speech-LLM Alignment. We first align the speech embeddings with the LLM input space by training on automatic speech recognition (ASR) task using approximately 90K hours of public speech data (VoxPopuli [30], MLS [31], Common Voice [32], People’s Speech [33]). During this stage, only the projection layer is updated, while all other model parameters remain frozen.

Conversation Pretraining. We then train on a mixture of publicly available multi-party conversation datasets (AMI [34], Fisher [35]). Since public corpora do not have a notion of “assistant” speaker, we simulate one by rotating the assistant role among all speakers except the conversation initiator. For a conversation with N participants, this yields $N - 1$ training instances.

Role-Conditioning Training. In the final training stage, we fine-tune exclusively on the *RolePlayConv* (Section 2.2) dataset, where the assistant’s role is specified via the system prompt.

For the latter two stages, we fine-tune the LLM parameters using low-rank adaptation [36], while keeping the speech encoder frozen throughout all stages (13.4M trainable parameters). We use the Adam optimizer with a two-phase learning rate schedule with the peak learning rate of 1×10^{-5} .

3.2. Evaluation Setup

Evaluation Datasets. We evaluate our models on two multi-party spoken conversation datasets. The first is NOTSOFAR-1 [37] (*NSF-I*), which consists of real recordings of formal meetings with approximately four speakers per session and an average duration of six minutes. As *NSF-I* lacks role labels, we designate one speaker as the “assistant” using a hybrid procedure that aggregates LLM-based rankings of assistant-like behavior with independent human evaluations. A role description is then generated for the speaker selected as assistant.

The second evaluation set is constructed following the same data generation pipeline as *RolePlayConv* (Section 2.2). To ensure a strict separation from training conditions, we use a zero-shot set of roles unseen during training and a different LLM, QwQ-32B [38], for the evaluation set construction.

Inference Configuration. During evaluation, each incoming audio chunk is treated as a potential turn-taking point. We adopt a dynamic chunking scheme, where chunk durations are uniformly sampled between 0.5 s and 3 s, and segmentation always occurs at speaker boundaries. As dynamic chunking introduces non-determinism, we report metrics averaged over 10 evaluation runs. In Section 4.2, we show that fixed-size chunking is not suitable for turn-taking evaluation. To obtain turn-taking predictions, we teacher-force the ground-truth conversation context together with the audio chunk under evaluation

before invoking generation. Greedy decoding is used by default. When a thinking control token (<think>) is emitted, we switch to sampling mode to generate the reasoning tokens. Similarly, upon emitting a turn-taking control token, sampling is used to generate the assistant’s response. Sampling is performed with *Temperature*=0.7, *TopP*=0.8, and *TopK*=20.

Evaluation Metrics. We evaluate predicted turn-taking decisions against ground-truth annotations at the chunk level using standard binary classification metrics: Precision, Recall, F1-score, Macro-averaged Accuracy, and False Positive Rate. In addition, we report the *Reactive Miss Rate*, which measures the fraction of missed turn-taking opportunities where a speaker directly addresses the assistant. To account for the inherent subjectivity of turn-taking and role adherence, we also conduct LLM-as-a-Judge evaluation [39] using *Claude-Sonnet-3.5* [40]. The judge model is provided with the assistant’s role description and the conversation history, and is prompted to assess (i) the appropriateness of the turn-taking decision with respect to the assigned role on a 0–1 scale, and (ii) role fidelity of the response content on a 1–10 scale. A small-scale human evaluation conducted on 100 instances shows strong agreement between human judgments and LLM-based scores, with a Spearman rank correlation of $\rho = 0.87$.

Baselines. We compare against two non-role-conditioned voice agent baselines that do not support role conditioning. The first is Moshi [5], designed for dyadic conversations (we use *Moshika*, the female-voice variant). Because Moshi operates at a frame rate of 12.5 Hz and produces continuous streaming outputs, we evaluate its turn-taking predictions using a tolerance window of ± 0.5 s around the ground-truth turn-taking points. The second baseline is the model fine-tuned on the *RolePlayConv* dataset without role-conditioning after the *Conversation Pretraining* stage (Section 3.1), which we refer to as *MP-Baseline*. As this model is trained on the same dataset as *ModeratorLM* models but without role-conditioning signals, it allows us to isolate the effect of role-conditioning on turn-taking behavior.

4. Results

4.1. Main Results

Table 2 compares *ModeratorLM* variants against non-role-conditioned baselines on *NSF-I* and *RolePlayConv* datasets. *Moshi*, trained on dyadic conversations, fails to generalize to multi-speaker settings, exhibiting very low recall and high false positive rates. The *MP-Baseline*, trained on multi-party conversations but without role conditioning, reduces interruptions but lacks the precision required for role-conditioned behavior. Both non-role-conditioned models exhibit substantially lower recall on *NSF-I*, which contains natural human-to-human conversations with frequent backchannels, interruptions, and overlapping speech. In contrast to the more structured user–assistant interactions in *RolePlayConv*, no participant consistently fol-

Model	Turn-Taking	Response
MP-Baseline	0.58	4.6
ModeratorLM	0.68	6.9
ModeratorLM-Think	0.72	7.4

Table 3: *LLM-as-a-judge evaluation of role fidelity for turn-taking ([0,1] scale) and responses ([0,10] scale) on RolePlayConv evaluation set.*

Roles:
<ul style="list-style-type: none"> Role A: A 46-year-old American news anchor with a crisp, authoritative delivery who enjoys live reporting, moderating panels, and investigative journalism. Role B: A 46-year-old American news anchor with a composed, non-intrusive demeanor who enjoys live reporting, moderating panels, and investigative journalism.
Conversation Context:
Nisha: Hey everyone, let's talk about leveraging influence and leadership.
Lance: Totally agree Nisha, it's crucial for driving change.
Thoughts:
<ul style="list-style-type: none"> Role A: Nisha initiated the topic, and Lance reinforced its importance. As a moderator in panels before, this discussion felt like an opening to set clear parameters for actionable insights without derailing their momentum into abstract theorizing. The first step is structure clarity through concise statements. Role B: Nisha initiated the topic, and Lance is affirming it. The conversation needs to build momentum before interjecting. Sometimes, silence allows ideas space without pressure from a passive anchor who values structured dialogue over interruption.

Table 4: *Reasoning traces generated by ModeratorLM-Think under two different roles for the same conversation context.*

lows assistant-style turn-taking patterns in *NSF-1*.

The *ModeratorLM* model variants perform substantially better across both datasets. *ModeratorLM* achieves high precision with very low false positives but remains conservative, missing some valid response opportunities. Incorporating chain-of-thought reasoning in *ModeratorLM-Think* improves this trade-off, yielding higher recall while maintaining low false positive rates, suggesting that explicit reasoning helps the model better interpret conversational context and role-specific obligations. Moreover, explicit reasoning reduces the miss rate at reactive turn-taking points.

Table 3 presents the LLM-as-a-Judge evaluation of role fidelity, where subjective trends closely mirror objective metrics. *ModeratorLM-Think* achieves the highest scores, indicating that the explicit reasoning stage serves as a foundational alignment layer. By first analyzing conversational context and role-specific obligations, the model ensures the decision of *when to speak* is cognitively grounded. This structured dependency yields turns that are both accurately timed and stylistically consistent. Finally, Table 4 provides qualitative examples demonstrating how the model adapts its reasoning under different roles given the same conversational context, resulting in distinct turn-taking decisions.

4.2. Ablation Studies

Table 5 reports results under different evaluation configurations designed to analyze the impact of chunking strategies and transcription availability on turn-taking. We compare the default *dynamic* chunking policy with two alternatives: (i) *Fixed*, where a constant chunk size is applied across all utterances; and (ii) *Turn-Fixed*, where the chunk size is constant within each

Setup	ModeratorLM			ModeratorLM-Th.		
	@P	@R	@A	@P	@R	@A
Default	0.71	0.57	0.76	0.79	0.82	0.91
No Transcription	0.42	0.14	0.57	0.39	0.42	0.57
ASR Hypotheses	0.68	0.56	0.76	0.75	0.80	0.90
GT Thoughts	–	–	–	0.95	0.95	0.97
Fixed (2 s)	0.88	0.78	0.88	0.82	0.82	0.91
Turn-Fixed	0.84	0.60	0.80	0.75	0.81	0.91

Table 5: *Evaluation results on RolePlayConv under different setups. @P, @R, and @A denote precision, recall, and macro-averaged accuracy, respectively. GT denotes ground truth. The default setup uses a fully dynamic chunking policy.*

speaker turn but may vary across turns, even within the same utterance. We define an *utterance* as a sequence of consecutive speaker turns occurring before the assistant responds, based on ground-truth annotations. For both alternative policies, segmentation is aligned with utterance boundaries, resulting in a variable-length final chunk for each utterance.

Under these settings, *ModeratorLM* exhibits higher turn-taking precision with both alternative chunking strategies compared to dynamic chunking. In the *Fixed* setting, recall also improves, as turn-taking chunks are typically shorter, while non-turn-taking chunks remain constant in length. This behavior suggests that *ModeratorLM* relies substantially on chunk-size cues when making turn-taking decisions. In contrast, *ModeratorLM-Think* is notably less sensitive to chunking strategy, instead relying on explicit reasoning traces. This robustness helps explain its near-perfect performance when ground-truth reasoning traces are provided. However, the apparent gains from fixed-size chunking are partly an artifact of the evaluation setup. In realistic inference scenarios, ground-truth utterance boundaries are unavailable. A fixed chunking policy must process an entire chunk before passing it to the model, which can delay or obscure turn-taking cues that emerge mid-chunk, leading to missed or mistimed interventions.

We also study the effect of transcriptions of user turns. Specifically, we evaluate two settings: (i) *No Transcription*, where no textual input is provided; and (ii) *ASR Hypotheses*, where transcripts are obtained using the streaming Kyutai-STT-2.6B model [41] independently on each user audio channel (word error rate of 6.7%). Performance degrades substantially in the *No Transcription* setting, indicating that the model relies heavily on textual information for turn-taking decisions. In contrast, using ASR hypotheses results in only minor degradation, demonstrating robustness to realistic transcription errors.

5. Conclusions

In this work, we introduced a role-playing voice agent for multi-party conversations that modulates turn-taking behavior based on an assigned role. Experimental results show that role-conditioned fine-tuning yields turn-taking decisions better aligned with configured preferences, and that incorporating chain-of-thought reasoning further improves role fidelity. Our analysis also highlights the importance of dynamic chunking during training; without it, models tend to overfit to chunk length rather than conversational context. Future work may explore full-duplex multi-party voice agents and turn-taking under simultaneous listening and speaking.

6. Acknowledgments

We would like to thank Ajay Srinivasamurthy, Volker Leutnant, Adam Kaplan, Andreas Schwarz, Raghavendra Bilgi and Sri Garimella for their support and valuable feedback.

7. References

- [1] X. Wang, Y. Li, C. Fu, Y. Shen, L. Xie, K. Li, X. Sun, and L. Ma, “Freeze-omni: A smart and low latency speech-to-speech dialogue model with frozen llm,” *arXiv preprint arXiv:2411.00774*, 2024.
- [2] Z. Xie and C. Wu, “Mini-omni: Language models can hear, talk while thinking in streaming, 2024,” *URL https://arxiv.org/abs/2408.16725*, 2024.
- [3] Q. Chen, Y. Chen, Y. Chen, M. Chen, Y. Chen, C. Deng, Z. Du, R. Gao, C. Gao, Z. Gao *et al.*, “Minmo: A multimodal large language model for seamless voice interaction,” *arXiv preprint arXiv:2501.06282*, 2025.
- [4] D. Ding, Z. Ju, Y. Leng, S. Liu, T. Liu, Z. Shang, K. Shen, W. Song, X. Tan, H. Tang *et al.*, “Kimi-audio technical report,” *arXiv preprint arXiv:2504.18425*, 2025.
- [5] A. Défossez, L. Mazaré, M. Orsini, A. Royer, P. Pérez, H. Jégou, E. Grave, and N. Zeghidour, “Moshi: a speech-text foundation model for real-time dialogue,” *arXiv preprint arXiv:2410.00037*, 2024.
- [6] Z. Ma, Y. Song, C. Du, J. Cong, Z. Chen, Y. Wang, Y. Wang, and X. Chen, “Language model can listen while speaking,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 23, 2025, pp. 24 831–24 839.
- [7] W. Yu, S. Wang, X. Yang, X. Chen, X. Tian, J. Zhang, G. Sun, L. Lu, Y. Wang, and C. Zhang, “Salmonn-omni: A codec-free llm for full-duplex speech understanding and generation,” *arXiv preprint arXiv:2411.18138*, 2024.
- [8] S. Arora, Z. Lu, C.-C. Chiu, R. Pang, and S. Watanabe, “Talking turns: Benchmarking audio foundation models on turn-taking dynamics,” *arXiv preprint arXiv:2503.01174*, 2025.
- [9] G.-T. Lin, S.-Y. S. Kuan, Q. Wang, J. Lian, T. Li, and H.-y. Lee, “Full-duplex-bench v1. 5: Evaluating overlap handling for full-duplex speech models,” *arXiv preprint arXiv:2507.23159*, 2025.
- [10] S. Sapkota, M. S. Hasan, M. Shah, and S. Karmaker, “Multi-party conversational agents: A survey,” *arXiv preprint arXiv:2505.18845*, 2025.
- [11] L. Zhou, J. Gao, D. Li, and H.-Y. Shum, “The design and implementation of xiaoice, an empathetic social chatbot,” *Computational Linguistics*, vol. 46, no. 1, pp. 53–93, 2020.
- [12] J. Chen, X. Wang, R. Xu, S. Yuan, Y. Zhang, W. Shi, J. Xie, S. Li, R. Yang, T. Zhu *et al.*, “From persona to personalization: A survey on role-playing language agents. arxiv [cs. cl](april 2024),” 2024.
- [13] N. Wang, Z. Peng, H. Que, J. Liu, W. Zhou, Y. Wu, H. Guo, R. Gan, Z. Ni, J. Yang *et al.*, “Rolellm: Benchmarking, eliciting, and enhancing role-playing abilities of large language models,” in *Findings of the Association for Computational Linguistics: ACL 2024*, 2024, pp. 14 743–14 777.
- [14] J. Zhou, Z. Chen, D. Wan, B. Wen, Y. Song, J. Yu, Y. Huang, P. Ke, G. Bi, L. Peng *et al.*, “Characterglm: Customizing social characters with large language models,” in *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, 2024, pp. 1457–1476.
- [15] Y. Ran, X. Wang, R. Xu, X. Yuan, J. Liang, Y. Xiao, and D. Yang, “Capturing minds, not just words: Enhancing role-playing language models with personality-indicative data,” in *Findings of the Association for Computational Linguistics: EMNLP 2024*, 2024, pp. 14 566–14 576.
- [16] R. Xu, X. Wang, J. Chen, S. Yuan, X. Yuan, J. Liang, Z. Chen, X. Dong, and Y. Xiao, “Character is destiny: Can role-playing language agents make persona-driven decisions?” *arXiv preprint arXiv:2404.12138*, 2024.
- [17] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou *et al.*, “Chain-of-thought prompting elicits reasoning in large language models,” *Advances in neural information processing systems*, vol. 35, pp. 24 824–24 837, 2022.
- [18] P. Pandey, R. V. Swaminathan, K. Girish, A. Sen, J. Xie, G. P. Strimel, and A. Schwarz, “Sift-50m: A large-scale multilingual dataset for speech instruction fine-tuning,” *arXiv preprint arXiv:2504.09081*, 2025.
- [19] J. Jia, G. Keren, W. Zhou, E. Lakomkin, X. Zhang, C. Wu, F. Seide, J. Mahadeokar, and O. Kalinli, “Efficient streaming llm for speech recognition,” in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025, pp. 1–5.
- [20] Q. Zhang, L. Cheng, C. Deng, Q. Chen, W. Wang, S. Zheng, J. Liu, H. Yu, C. Tan, Z. Du *et al.*, “Omniflatten: An end-to-end gpt model for seamless voice conversation,” *arXiv preprint arXiv:2410.17799*, 2024.
- [21] J. Xu, Z. Guo, J. He, H. Hu, T. He, S. Bai, K. Chen, J. Wang, Y. Fan, K. Dang *et al.*, “Qwen2. 5-omni technical report,” *arXiv preprint arXiv:2503.20215*, 2025.
- [22] X. Wang, H. Wang, Y. Zhang, X. Yuan, R. Xu, J.-t. Huang, S. Yuan, H. Guo, J. Chen, S. Zhou *et al.*, “Coser: Coordinating llm-based persona simulation of established roles,” in *Forty-second International Conference on Machine Learning*, 2025.
- [23] S. Poria, D. Hazarika, N. Majumder, G. Naik, E. Cambria, and R. Mihalcea, “Meld: A multimodal multi-party dataset for emotion recognition in conversations,” in *Proceedings of the 57th annual meeting of the association for computational linguistics*, 2019, pp. 527–536.
- [24] D. Fu, X. Cheng, L. Li, X. Yang, L. Yang, and T. Jin, “Pachat: Persona-aware speech assistant for multi-party dialogue,” in *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 2025, pp. 29 313–29 330.
- [25] A. A. G. Intelligence, “The amazon nova family of models: Technical report and model card,” 2024.
- [26] Zypfra, “Zonos: an open-weight multilingual text-to-speech model (zonos-v0.1),” <https://github.com/Zypfra/Zonos>, 2025.
- [27] A. Yang, A. Li, B. Yang, B. Zhang, B. Hui, B. Zheng, B. Yu, C. Gao, C. Huang, C. Lv *et al.*, “Qwen3 technical report,” *arXiv preprint arXiv:2505.09388*, 2025.
- [28] P. Swietojanski, S. Braun, D. Can, T. F. Da Silva, A. Ghoshal, T. Hori, R. Hsiao, H. Mason, E. McDermott, H. Silovsky *et al.*, “Variable attention masking for configurable transformer transducer speech recognition,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [29] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, “Montreal forced aligner: Trainable text-speech alignment using kaldia,” in *Interspeech*, vol. 2017, 2017, pp. 498–502.
- [30] C. Wang, M. Riviere, A. Lee, A. Wu, C. Talnikar, D. Haziza, M. Williamson, J. Pino, and E. Dupoux, “Voxpopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation,” in *ACL*, 2021.
- [31] V. Pratap, Q. Xu, S. Aniceto *et al.*, “MIs: A large-scale multilingual dataset for scalable speech research,” in *Interspeech*, 2020.
- [32] R. Ardila, M. Branson, K. Davis *et al.*, “Common voice: A massively-multilingual speech corpus,” in *LREC*, 2020.
- [33] D. Galvez *et al.*, “The people’s speech: A large-scale diverse english speech recognition dataset for commercial usage,” *arXiv preprint arXiv:2111.09340*, 2021.
- [34] J. Carletta, S. Ashby, S. Bourban, M. Flynn, M. Guillemot, T. Hain, J. Kadlec, V. Karaiskos, W. Kraaij, M. Kronenthal *et al.*, “The ami meeting corpus: A pre-announcement,” in *International workshop on machine learning for multimodal interaction*. Springer, 2005, pp. 28–39.
- [35] C. Cieri, D. Miller, and K. Walker, “The fisher corpus: A resource for the next generations of speech-to-text.”

- [36] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, “Lora: Low-rank adaptation of large language models,” *arXiv preprint arXiv:2106.09685*, 2021.
- [37] A. Vinnikov, A. I. Mark, A. Hurvitz, I. Abramovski, S. Koubi, I. Gurvich, S. Peer, X. Xiao, B. Elizalde, N. Kanda *et al.*, “Notsofar-1 challenge: New datasets, baseline, and tasks for distant meeting transcription,” in *Proc. CHiME 2024*, 2024.
- [38] Q. Team, “Qwq-32b: Embracing the power of reinforcement learning,” 2025.
- [39] L. Zheng, W.-L. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E. Xing *et al.*, “Judging llm-as-a-judge with mt-bench and chatbot arena,” *Advances in neural information processing systems*, vol. 36, pp. 46 595–46 623, 2023.
- [40] Anthropic, “Claude 3.5 Sonnet,” <https://www.anthropic.com/news/claude-3-5-sonnet>, 2024.
- [41] N. Zeghidour, E. Kharitonov, M. Orsini, V. Volhejn, G. de Mermiesse, E. Grave, P. Pérez, L. Mazaré, and A. Défossez, “Streaming sequence-to-sequence learning with delayed streams modeling,” *arXiv preprint arXiv:2509.08753*, 2025.