

# AMPS: ASR with Multimodal Paraphrase Supervision

Abhishek Gupta\* Amruta Parulekar\* Sameep Chattopadhyay Preethi Jyothi  
Indian Institute of Technology Bombay, Mumbai, India

{abhishekumgupta, amrutaparulekar.iitb, sameep.ch.2002}@gmail.com, pjyothi@cse.iitb.ac.in

## Abstract

Spontaneous or conversational multilingual speech presents many challenges for state-of-the-art automatic speech recognition (ASR) systems. In this work, we present a new technique AMPS that augments a multilingual multimodal ASR system with paraphrase-based supervision for improved conversational ASR in multiple languages, including Hindi, Marathi, Malayalam, Kannada, and Nyanja. We use paraphrases of the reference transcriptions as additional supervision while training the multimodal ASR model and selectively invoke this paraphrase objective for utterances with poor ASR performance. Using AMPS with a state-of-the-art multimodal model SeamlessM4T, we obtain significant relative reductions in word error rates (WERs) of up to 5%. We present detailed analyses of our system using both objective and human evaluation metrics.

## 1 Introduction

Automatic speech recognition (ASR) systems have shown considerable progress in recent years but still falter when subjected to spontaneous conversational speech containing disfluencies, loosely articulated sounds, and other noise factors (Gabler et al., 2023). This degradation in ASR performance could be largely attributed to the unavailability of labeled spontaneous speech in most languages. How can we effectively utilize the limited quantities of existing labeled spontaneous speech? Towards this, we propose AMPS (ASR with Multimodal Paraphrase Supervision) that augments an existing multilingual multimodal ASR system with paraphrase-based supervision to improve ASR performance on spontaneous speech in multiple languages.

Unlike standalone ASR models that are exclusively trained to perform ASR, multimodal models (such as SpeechT5 (Ao et al., 2022), MAESTRO (Chen et al., 2022), etc.) are trained on multiple

tasks *including* ASR using speech and text data in various paired (and unpaired) forms. We focus on one such multilingual multimodal model, SeamlessM4T (Communication et al., 2023), that consists of dual encoders for speech and text and a shared text decoder, thus creating both speech-to-text and text-to-text pathways.

AMPS<sup>1</sup> leverages the multimodal nature of SeamlessM4T by introducing a paraphrasing objective jointly with ASR. Along with using spontaneous speech and its corresponding transcription to train the speech-to-text pathway in SeamlessM4T, AMPS also uses paraphrases of the reference transcriptions as additional supervision to train the text-to-text pathway. We selectively employ paraphrase-based augmentation during training when the ASR loss is high (as determined by a predetermined threshold); high ASR loss is typically triggered by noise or poorly enunciated words in spontaneous speech. This selective intervention offers the model an alternate path of opting for semantically close words and phrases when the audio is not very clear. It is important that the paraphrases should not significantly differ in word order from the original transcripts, thus enabling the model to easily align representations of speech, text, and its paraphrase.

With AMPS, we derive significant improvements in ASR for spontaneous speech in Hindi, Marathi, Malayalam, Kannada, and Nyanja compared to strong ASR-only finetuned baselines. We report improvements not only in terms of word error rate (WER) reductions but also using semantic evaluation metrics. We also conduct a detailed human evaluation comparing the outputs of AMPS with the outputs from finetuning only with the ASR objective and show consistent improvements in human scores. We also present many ablations, including different paraphrasing techniques, the influence of

\*These authors contributed equally to this work.

<sup>1</sup>Code for AMPS is available at <https://github.com/csalt-research/amps-asr>.

varying thresholds on the performance of AMPS, and using varying amounts of training data. We envision that techniques like AMPS could be used to improve ASR of atypical speech for people with speech impairments where comprehensibility of the transcripts is critical (more than faithfulness of transcripts to the underlying speech, as highlighted in very recent work by Tomanek et al. (2024)).

## 2 Related Work

In recent years, multimodal models for speech recognition have gained significant recognition (Ao et al., 2022; Chen et al., 2022; Rubenstein et al., 2023; Zhang et al., 2023). These models are capable of processing both speech and text inputs and can be adapted for tasks such as translation and speech generation. A notable example is Meta AI’s SeamlessM4T (Communication et al., 2023), which can support nearly 100 languages. One of the key advantages of such models is their ability to exploit text-only training to fine-tune shared parameters in the ASR pipeline. Some of the recent work on text-based adaptation for ASR models include Vuong et al. (2023); Bataev et al. (2023); Chen et al. (2023); Mittal et al. (2023). One potential approach for leveraging text-only data for ASR finetuning is through training the text decoder with a paraphrasing objective. Emerging research (Yu et al., 2023) has shown that text paraphrasing can be used to augment LLM performance but we are the first to show how paraphrases can be used to improve ASR. Tomanek et al. (2024) is a recent study focusing on meaning preservation in disordered speech transcription, but do not offer any technique to help improve meaning preservation in ASR outputs.

## 3 Methodology

AMPS scaffolds on a multimodal base model comprising a speech encoder, a text encoder, and a shared decoder that takes inputs from both encoders. SeamlessM4T is an example of such a model, capable of performing multiple tasks including text-to-text translation (T2T), and speech-to-text transcription/translation (S2T). We introduce a new auxiliary task of text-to-text paraphrasing. This allows the model to predict words that are semantically similar and fit within the context of the sentence, without significantly altering its word order. The shared decoder architecture of SeamlessM4T allows us to exploit common parameters of both S2T and T2T pipelines and enhance the ASR performance of the model.

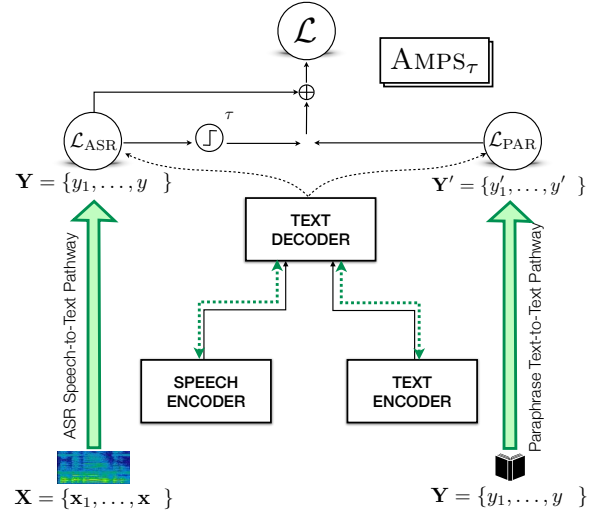


Figure 1: **Multimodal AMPS<sub>τ</sub> Pipeline.** AMPS<sub>τ</sub> applies a dual pass through the S2T pipeline with an ASR objective and the T2T pipeline with a paraphrasing objective. The paraphrasing loss is only incorporated when the ASR loss exceeds a predefined threshold.

Formally, consider a speech utterance  $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_L \mid \mathbf{x}_i \in \mathbb{R}^d\}$  with its corresponding transcript  $\mathbf{Y} = \{y_1, y_2, \dots, y_N\}$ . For a transcript  $\mathbf{Y}$ , we generate a paraphrase  $\mathbf{Y}' = \{y'_1, y'_2, \dots, y'_M\}$ . Given a labeled instance  $\{\mathbf{X}, \mathbf{Y}, \mathbf{Y}'\}$ , the ASR, paraphrase, and the AMPS loss functions are as follows.

$$\begin{aligned} \mathcal{L}_{\text{ASR}} &= \sum_{t=1}^N \log p_{\theta}(y_t \mid y_{<t}, \mathbf{X}), \\ \mathcal{L}_{\text{PAR}} &= \sum_{t=1}^M \log p_{\phi}(y'_t \mid y'_{<t}, \mathbf{Y}), \\ \mathcal{L}_{\text{AMPS}} &= \mathcal{L}_{\text{ASR}} + \mathcal{L}_{\text{PAR}}. \end{aligned}$$

For each batch, we pass the audio through the S2T pathway and compute the ASR loss between the predicted and ground-truth transcripts. We also pass the ground-truth transcripts as input through the T2T pathway with paraphrase-based supervision to compute  $\mathcal{L}_{\text{PAR}}$ . Figure 1 illustrates a schematic of our proposed architecture.

**AMPS<sub>τ</sub>: Loss Function Thresholding.** We aim at improving the model’s performance in noisy regions where the ASR loss is high by selectively triggering the paraphrase objective only when the ASR loss exceeds a predefined threshold  $\tau$ .

Thus, the loss for the system is given by

$$\mathcal{L}_{\text{AMPS}_{\tau}} = \begin{cases} \mathcal{L}_{\text{ASR}} + \mathcal{L}_{\text{PAR}} & \text{if } \mathcal{L}_{\text{ASR}} > \tau, \\ \mathcal{L}_{\text{ASR}} & \text{otherwise,} \end{cases} \quad (1)$$

Language	Evaluation Type	Direct Inference	All Data			Hard 100		$\Delta = \text{AMPS}_\tau - \text{ASR}$	
	Configuration	-	ASR	AMPS	$\text{AMPS}_\tau$	ASR	$\text{AMPS}_\tau$	$\Delta\text{Hard}$	$\Delta\text{All}$
Marathi	WER ↓	38.65	21.18	21.58	<b>20.20</b>	48.91	42.79	-6.12	-0.98
	METEOR ↑	59.84	73.32	<b>77.67</b>	76.62	54.13	58.45	4.32	3.30
	BERTScore ↑	81.01	90.40	<b>92.31</b>	91.92	84.73	85.82	0.99	1.52
Hindi	WER ↓	29.16	20.63	20.83	<b>20.12</b>	49.09	45.91	-3.18	-0.51
	METEOR ↑	72.25	81.04	81.38	<b>81.56</b>	57.66	60.91	3.25	0.52
	BERTScore ↑	88.55	93.60	93.65	<b>93.76</b>	84.46	85.44	0.98	0.16
Malayalam	WER ↓	56.15	42.06	42.09	<b>39.97</b>	74.86	64.66	-10.2	-2.09
	METEOR ↑	43.69	60.39	60.31	<b>62.01</b>	32.48	40.58	8.10	1.62
	BERTScore ↑	84.35	91.50	91.56	<b>92.02</b>	85.40	87.41	2.01	0.52
Kannada	WER ↓	69.29	41.41	40.10	<b>39.50</b>	72.23	67.58	-4.65	-1.91
	METEOR ↑	31.13	60.84	61.27	<b>61.68</b>	33.44	38.30	4.86	0.84
	BERTScore ↑	76.65	89.84	90.21	<b>90.41</b>	82.36	85.54	3.18	0.57

Table 1: Comparing the performance of pure ASR, AMPS, and  $\text{AMPS}_\tau$  systems using 50 hours of training data with round-trip translated paraphrases. Best overall scores for each metric are highlighted in  .

where  $\tau$  is a hyperparameter chosen based on ASR validation losses. Henceforth, AMPS with the best threshold will be referred to as  $\text{AMPS}_\tau$ .  $\tau$  values for various experiments are in Appendix A.

## 4 Experimental Setup

For all our experiments, we use the SeamlessM4T multilingual multimodal model (Communication et al., 2023). The text encoder and decoder modules are initialized using Meta’s No Language Left Behind (NLLB) model (Team et al., 2022). The speech encoder in SeamlessM4T uses Wav2Vec-BERT 2.0 (Kessler et al., 2021), which is trained on over a million hours of unlabeled speech data. Further model details are in Appendix B.1.

**Datasets.** The IndicVoices dataset (Javed et al., 2024b) is a large collection of natural speech (74% extempore, 17% conversational and 9% read) in 22 Indic languages. Among the languages we chose, Marathi, Kannada, and Malayalam are classified as low-resource by SeamlessM4T (Communication et al., 2023), while Hindi is medium-resource. IndicVoices is the only multilingual open-source Indian speech corpus containing spontaneous speech and amongst the very few sources published after SeamlessM4T’s release.<sup>2</sup> We also performed experiments on Nyanja (a low-resource language from Zambia) from the Zambezi-Voice dataset (Sikasote et al., 2023).

We use roughly 50 hours of (predominantly conversational, henceforth referred to as *mixed*) training data for each of the four Indian languages. For

Hindi, we also simulate a very low-resource setting with random 5-hour samples of mixed and read training speech. For Nyanja, we used 5 hours of training data. (For Indic languages, our test sets are the validation sets that are part of IndicVoices. For Nyanja, we use the existing test set.) Given the limited amount of training data, we use parameter-efficient finetuning of adapter layers (Houlsby et al., 2019) in the speech encoder and text decoder layers of the SeamlessM4T model; more implementation details are in Appendix B.2.

**Paraphrasing.** We translated the reference transcriptions into English using IndicTrans-2 (Gala et al., 2023) for the Indic languages and NLLB (Team et al., 2022) for Nyanja before translating them back to their original languages. For the Hindi mixed 5-hr setting, we experimented with top- $K$ ,  $K = 50$ , and nucleus (top- $P$ ,  $P = 0.95$ ) sampling during round-trip translation to produce more diverse paraphrases. We also explored generating paraphrases using the multilingual LLM Aya-23 (Üstün et al., 2024). The exact prompt and other details are in Appendix C and D.2. We used round-trip translation-based paraphrases for all the 50-hour experiments due to poor-quality LLM paraphrases for low-resource languages like Malayalam.

**Evaluation Metrics.** Evaluation metrics used were Word Error Rate (WER), METEOR and the F1 score provided by BERTScore. More details are provided in Appendix E.

## 5 Experiments and Results

Table 1 shows the main results for all the 50-hour Indian-language experiments.  $\text{AMPS}_\tau$  consistently

<sup>2</sup>This dataset was chosen also to ensure that there was no data leakage between the SeamlessM4T training data and the evaluation sets.

Language	Paraphrase Type	Direct Inference	Read Speech			Mixed Speech						
	Configuration	-	RT Trans			RT Trans			LLM-Para		TK+Nuc RT Trans	
			ASR	AMPS	AMPS <sub>τ</sub>	ASR	AMPS	AMPS <sub>τ</sub>	AMPS	AMPS <sub>τ</sub>	AMPS	AMPS <sub>τ</sub>
Hindi	WER ↓	29.16	<b>28.19</b>	28.94	28.57	23.14	23.14	<b>22.80</b>	22.35	<b>22.20</b>	<b>22.58</b>	22.81
	METEOR ↑	72.25	<b>74.36</b>	73.58	73.91	79.10	78.86	<b>78.93</b>	79.25	<b>79.28</b>	<b>79.27</b>	79.11
	BERTScore ↑	88.55	<b>90.39</b>	89.86	90.13	92.60	92.59	<b>92.78</b>	<b>92.89</b>	92.90	<b>92.63</b>	92.62

Table 2: Comparing ASR, AMPS and AMPS<sub>τ</sub> systems using 5 hours of mixed (conversational and read) speech with round-trip translations (RT Trans), LLM paraphrasing and top-K + nucleus paraphrasing.

Language	ASR	AMPS	AMPS <sub>τ</sub>
Marathi	4.199	4.271	<b>4.314</b>
Hindi	3.608	3.625	<b>3.689</b>
Malayalam	3.635	3.688	<b>3.902</b>
Kannada	3.433	3.542	<b>3.597</b>

Table 3: Comparison of human annotation results for ASR, AMPS and AMPS<sub>τ</sub> on a scale from 0 to 5.

performs best compared to ASR, and the WER reductions are statistically significant (at  $p < 0.05$  using the mapsswe test).<sup>3</sup> Apart from the overall scores in *All Data*, we sorted the transcriptions in descending order of WER using pure ASR and averaged metrics were calculated for both pure ASR and AMPS<sub>τ</sub> for the first 100 (hardest) sentences. Improvements from ASR to AMPS<sub>τ</sub> for these hardest 100 predictions are labeled  $\Delta Hard$  in Table 1. We see that  $\Delta Hard$  consistently exceeds  $\Delta All$ , indicating that the most improvement is observed in cases where pure ASR performs poorly. This supports the thresholding approach that triggers the paraphrase loss only when pure ASR predictions fall below a threshold. From our manual inspection of Hindi samples in the hardest-100 subset, we observe examples where pure ASR tends to produce acoustically similar but incorrect words, while AMPS<sub>τ</sub> correctly identifies the words. For example, pure ASR misrecognized “hua” (meaning ‘is’) as “ugwa” (meaning ‘grows’) in a Hindi example; AMPS<sub>τ</sub> gets this example right.

## 5.1 Comparing Paraphrase Techniques

Table 2 shows results from training on 5 hrs of read/mixed Hindi speech and different paraphras-

<sup>3</sup>We also trained a variant where instances with a ASR loss were downweighted and instances with a high ASR loss were upweighted, thus forcing the model to focus more on the latter. This performed comparably to our baseline ASR model.

ing techniques with mixed speech. Here, by mixed speech, we refer to a mixture of both read and conversational speech. Unsurprisingly, training on mixed speech yields significantly lower WERs compared to training on read speech. The highest performance gains were obtained using LLM paraphrasing for Hindi, suggesting that the LLM is a good option for medium-resource languages like Hindi. LLM outputs are subpar for low-resource languages like Kannada, and hence are not an option. Comprehensive results comparing the paraphrase techniques for other languages are given in Appendix F and G.

## 5.2 Human Evaluation

The transcription capabilities of ASR, AMPS, and AMPS<sub>τ</sub> models were verified through extensive human evaluation of the utterances with differing model outputs. The annotators reviewed 172, 153, 216, and 229 instances for Hindi, Marathi, Kannada, and Malayalam, respectively, giving a max score of 5 for a perfect transcript and penalizing them for minor errors (spellings, etc.) and major errors (incorrect semantics). The annotators were asked not to penalize a semantically identical word that differs from the speech. More details and scoring guidelines are provided in Appendix H and qualitative examples are in Appendix D.1. Table 3 shows the averaged scores with AMPS<sub>τ</sub> consistently performing the best across all languages.

## 5.3 AMPS for Nyanja

Table 4 shows overall results<sup>4</sup> on Nyanja with 5 hours of training data and round-trip translated paraphrases. Again, AMPS<sub>τ</sub> performs the best, showing that AMPS could be applied to diverse languages across language families.

<sup>4</sup>Only WER and METEOR are reported. BERTScore does not support Nyanja.



Language	Config.	Direct Inference	ASR	AMPS	AMPS <sub><math>\tau</math></sub>
Nyanja	WER ↓	42.34	22.16	21.90	<b>21.59</b>
	METEOR ↑	66.71	79.25	79.30	<b>80.10</b>

Table 4: Comparison of WER (%) and METEOR for ASR, AMPS and AMPS <sub>$\tau$</sub>  for 5 hours Nyanja speech with round-trip translated paraphrases.

## 5.4 Conclusion

This work introduces a novel paraphrase-based supervision technique AMPS to improve the ASR performance of spontaneous speech in multimodal models. This auxiliary supervision makes the model more robust and helps the model generalize better, especially in utterances with large ASR errors. We show significant ASR improvements on multiple and diverse languages and further validate these improvements via a thorough human evaluation.

The broader idea of using textual supervision, as we did with paraphrases, to improve speech understanding is an interesting avenue to explore further. Future work will investigate how techniques like AMPS could be used to improve ASR for atypical speech. Also, we used a predefined threshold on the ASR loss to trigger the paraphrase objective; this could be made a learnable quantity.

## 6 Acknowledgements

The authors thank the anonymous reviewers for their constructive feedback that improved the quality of the draft. The last author gratefully acknowledges support from the Amazon IITB AI ML Initiative.

## Limitations

The primary limitation of our study was the lack of any appropriate pre-existing evaluation metric for the task. When supervising with paraphrases, the model often predicts semantically similar words or phrases that do not exactly match the transcript, making traditional metrics like Word Error Rate (WER) overly harsh for such cases. While BERTScore addresses semantic similarity, recent research suggests using LLMs to directly assess whether sentence meaning is preserved (Tomanek et al., 2024). In the future, we plan to adopt LLM-based evaluation alongside human reviews to improve assessment.

A second limitation was the occurrence of transliterated English words caused minor spelling

errors in the model. We plan to mitigate this in the future by introducing code-switched words in our paraphrases to teach the model to associate the transliterated English words with their Latin script counterparts. Multilingual models like SeamlessM4T possess the unique ability to link semantically similar words across languages, thus comprehending code-switched speech easily and we aim to leverage this ability as future work.

Additionally, the threshold value  $\tau$  is manually defined and not a dynamic value that is learned across languages. In future work, we plan to make this threshold a learnable parameter.

## References

- Junyi Ao, Rui Wang, Long Zhou, Chengyi Wang, Shuo Ren, Yu Wu, Shujie Liu, Tom Ko, Qing Li, Yu Zhang, Zhihua Wei, Yao Qian, Jinyu Li, and Furu Wei. 2022. [SpeechT5: Unified-modal encoder-decoder pre-training for spoken language processing](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5723–5738, Dublin, Ireland. Association for Computational Linguistics.
- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Vladimir Bataev, Roman Korostik, Evgeny Shabalin, Vitaly Lavrukhin, and Boris Ginsburg. 2023. [Text-only domain adaptation for end-to-end asr using integrated text-to-mel-spectrogram generator](#). In *Interspeech*.
- Chang Chen, Xun Gong, and Yanmin Qian. 2023. [Efficient text-only domain adaptation for ctc-based asr](#). In *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 1–7.
- Zhehuai Chen, Yu Zhang, Andrew Rosenberg, Bhuvana Ramabhadran, Pedro J. Moreno, Ankur Bapna, and Heiga Zen. 2022. [Maestro: Matched speech text representations through modality matching](#). In *Interspeech*.
- Seamless Communication, Loïc Barrault, Yu-An Chung, Mariano Cora Meglioli, David Dale, Ning Dong, Paul-Ambroise Duquenne, Hady Elsahar, Hongyu Gong, Kevin Heffernan, John Hoffman, Christopher Klaiber, Pengwei Li, Daniel Licht, Jean Maillard, Alice Rakotoarison, Kaushik Ram Sadagopan, Guillaume Wenzek, Ethan Ye, Bapi Akula, Peng-Jen Chen, Naji El Hachem, Brian Ellis, Gabriel Mejia Gonzalez, Justin Haaheim, Prangthip Hansanti, Russ

- Howes, Bernie Huang, Min-Jae Hwang, Hirofumi Inaguma, Somya Jain, Elahe Kalbassi, Amanda Kallet, Ilia Kulikov, Janice Lam, Daniel Li, Xutai Ma, Ruslan Mavlyutov, Benjamin Peloquin, Mohamed Ramadan, Abinash Ramakrishnan, Anna Sun, Kevin Tran, Tuan Tran, Igor Tufanov, Vish Vogeti, Carleigh Wood, Yilin Yang, Bokai Yu, Pierre Andrews, Can Balioglu, Marta R. Costa-jussà, Onur Celebi, Maha Elbayad, Cynthia Gao, Francisco Guzmán, Justine Kao, Ann Lee, Alexandre Mourachko, Juan Pino, Sravya Popuri, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, Paden Tomasello, Changhan Wang, Jeff Wang, and Skyler Wang. 2023. [Seamless4t: Massively multilingual & multimodal machine translation](#). *Preprint*, arXiv:2308.11596.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). *Preprint*, arXiv:1810.04805.
- Philipp Gabler, Bernhard C Geiger, Barbara Schuppler, and Roman Kern. 2023. Reconsidering read and spontaneous speech: Causal perspectives on the generation of training data for automatic speech recognition. *Information*, 14(2):137.
- Jay Gala, Pranjal A Chitale, A K Raghavan, Varun Gumma, Sumanth Doddapaneni, Aswanth Kumar M, Janki Atul Nawale, Anupama Sujatha, Ratish Pudupully, Vivek Raghavan, Pratyush Kumar, Mitesh M Khapra, Raj Dabre, and Anoop Kunchukuttan. 2023. [Indictrans2: Towards high-quality and accessible machine translation models for all 22 scheduled indian languages](#). *Transactions on Machine Learning Research*.
- Dan Hendrycks and Kevin Gimpel. 2023. [Gaussian error linear units \(gelus\)](#). *Preprint*, arXiv:1606.08415.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for nlp.
- Tahir Javed, Janki Nawale, Eldho George, Sakshi Joshi, Kaushal Bhogale, Deovrat Mehendale, Ishvinder Sethi, Aparna Ananthanarayanan, Hafsa Faquih, Pratiti Palit, Sneha Ravishankar, Saranya Sukumaran, Tripura Panchagnula, Sunjay Murali, Kunal Gandhi, Ambujavalli R, Manickam M, C Vijayanthi, Krishnan Karunganni, Pratyush Kumar, and Mitesh Khapra. 2024a. [IndicVoices: Towards building an inclusive multilingual speech dataset for Indian languages](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 10740–10782, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Tahir Javed, Janki Atul Nawale, Eldho Ittan George, Sakshi Joshi, Kaushal Santosh Bhogale, Deovrat Mehendale, Ishvinder Virender Sethi, Aparna Ananthanarayanan, Hafsa Faquih, Pratiti Palit, Sneha Ravishankar, Saranya Sukumaran, Tripura Panchagnula, Sunjay Murali, Kunal Sharad Gandhi, Ambujavalli R, Manickam K M, C Venkata Vijayanthi, Krishnan Srinivasa Raghavan Karunganni, Pratyush Kumar, and Mitesh M Khapra. 2024b. [Indicvoices: Towards building an inclusive multilingual speech dataset for indian languages](#). *Preprint*, arXiv:2403.01926.
- Divyanshu Kakwani, Anoop Kunchukuttan, Satish Golla, Gokul N.C., Avik Bhattacharyya, Mitesh M. Khapra, and Pratyush Kumar. 2020. IndicNLP Suite: Monolingual Corpora, Evaluation Benchmarks and Pre-trained Multilingual Language Models for Indian Languages. In *Findings of EMNLP*.
- Samuel Kessler, Bethan Thomas, and Salah Karout. 2021. [Continual-wav2vec2: an application of continual learning for self-supervised automatic speech recognition](#). *ArXiv*, abs/2107.13530.
- Ashish Mittal, Sunita Sarawagi, and Preethi Jyothi. 2023. [In-situ text-only adaptation of speech models with low-overhead speech imputations](#). In *The Eleventh International Conference on Learning Representations*.
- Omkar Patil, Rahul Singh, and Tarun Joshi. 2022. [Understanding metrics for paraphrasing](#). *ArXiv*, abs/2205.13119.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. [Robust speech recognition via large-scale weak supervision](#). *arXiv preprint*.
- Paul K. Rubenstein, Chulayuth Asawaroengchai, Duc Dung Nguyen, Ankur Bapna, Zalán Borsos, Félix de Chaumont Quitry, Peter Chen, Dalia El Badawy, Wei Han, Eugene Kharitonov, Hannah Muckenhirn, Dirk Padfield, James Qin, Danny Rozenberg, Tara Sainath, Johan Schalkwyk, Matt Sharifi, Michelle Tadmor Ramanovich, Marco Tagliasacchi, Alexandru Tudor, Mihajlo Velimirović, Damien Vincent, Jiahui Yu, Yongqiang Wang, Vicky Zayats, Neil Zeghidour, Yu Zhang, Zhishuai Zhang, Lukas Zilka, and Christian Frank. 2023. [Audiopalm: A large language model that can speak and listen](#). *Preprint*, arXiv:2306.12925.
- Lingfeng Shen, Haiyun Jiang, Lemao Liu, and Shuming Shi. 2022a. [Revisiting the evaluation metrics of paraphrase generation](#). *ArXiv*, abs/2202.08479.
- Lingfeng Shen, Lemao Liu, Haiyun Jiang, and Shuming Shi. 2022b. [On the evaluation metrics for paraphrase generation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3178–3190, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Claytone Sikasote, Kalinda Siaminwe, Stanly Mwape, Bangiwe Zulu, Mofya Phiri, Martin Phiri, David Zulu, Mayumbo Nyirenda, and Antonios Anastasopoulos. 2023. [Zambezi Voice: A Multilingual Speech Corpus for Zambian Languages](#). In *Proc. INTERSPEECH 2023*, pages 3984–3988.

NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Hefernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. [No language left behind: Scaling human-centered machine translation](#). *Preprint*, arXiv:2207.04672.

Katrin Tomanek, Jimmy Tobin, Subhashini Venugopalan, Richard Cave, Katie Seaver, Jordan R. Green, and Rus Heywood. 2024. [Large language models as a proxy for human evaluation in assessing the comprehensibility of disordered speech transcription](#). In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 10846–10850.

A Vaswani. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.

Tyler Vuong, Karel Mundnich, Dhanush Bekal, Veera Elluru, Srikanth Ronanki, and Sravan Bodapati. 2023. [AdaBERT-CTC: Leveraging BERT-CTC for text-only domain adaptation in ASR](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 364–371, Singapore. Association for Computational Linguistics.

Yijiong Yu, Yongfeng Huang, Zhixiao Qi, and Zhe Zhou. 2023. [Training with "paraphrasing the original text" improves long-context performance](#).

Dong Zhang, Shimin Li, Xin Zhang, Jun Zhan, Pengyu Wang, Yaqian Zhou, and Xipeng Qiu. 2023. [Speechgpt: Empowering large language models with intrinsic cross-modal conversational abilities](#). *Preprint*, arXiv:2305.11000.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). *Preprint*, arXiv:1904.09675.

Ahmet Üstün, Viraat Aryabumi, Zheng-Xin Yong, Wei-Yin Ko, Daniel D’souza, Gbemileke Onilude, Neel Bhandari, Shivalika Singh, Hui-Lee Ooi, Amr Kayid, Freddie Vargus, Phil Blunsom, Shayne Longpre, Niklas Muennighoff, Marzieh Fadaee, Julia Kreutzer, and Sara Hooker. 2024. [Aya model: An instruction finetuned open-access multilingual language model](#). *Preprint*, arXiv:2402.07827.

## Appendix

### A Thresholds for AMPS<sub>τ</sub>

Table 5 contains the iteratively obtained best thresholds for the training sets for our experiments. In case of inconsistency between different metrics, the best threshold was chosen using the validation WER for the pure ASR system.

Language	Read BT	Mixed BT	Mixed BT	Mixed LLM	Mixed Top-K BT
Hours	<5	50	5	5	5
Marathi	3.5	3.8	3.6	-	3.6
Hindi	3.2	3.2	3.6	3.6	3.6
Malayalam	3.8	3.8	3.4	-	3.4
Kannada	3.8	3.6	3.4	-	3.2
Nyanja	-	-	3.8	-	-

Table 5: Iteratively obtained threshold values for all the experimental datasets for AMPS<sub>τ</sub>.

### B AMPS for SeamlessM4T

For all our experiments, we used the SeamlessM4T medium model along with IndicVoices (Javed et al., 2024a), and Zambezi-voice (Sikasote et al., 2023) datasets. Both the data and the models are free and open-sourced.

#### B.1 Adapting SeamlessM4T

The SeamlessM4T (Medium) consists of 1.2B parameters. Full fine-tuning of these components using limited amounts of labeled data for low-resource languages may result in overfitting and degradation of ASR performance. To address these issues, parameter-efficient fine-tuning methods, such as the adapter framework, have become widely adopted in natural language processing tasks. Adapters have proven effective in low-resource ASR tasks, including accent and cross-lingual adaptation.

Formally, the operations performed in the  $i^{\text{th}}$  speech encoder layer can be described as follows:

$$\mathbf{H} = \text{MHA}(\mathbf{h}^{i-1}, \mathbf{h}^{i-1}, \mathbf{h}^{i-1})$$

$$\mathbf{C} = \text{Convolution}(\mathbf{H})$$

$$\hat{\mathbf{h}}^i = \text{FFN}(\mathbf{C})$$

$$\mathbf{h}^i = \text{Adapter}(\hat{\mathbf{h}}^i)$$

Similarly, the operations in the  $i^{\text{th}}$  decoder layer can be summarized as:

$$\begin{aligned}
\mathbf{D} &= \text{MHA}(\mathbf{d}^{i-1}, \mathbf{d}^{i-1}, \mathbf{d}^{i-1}) \\
\hat{\mathbf{D}} &= \text{MHA}(\mathbf{d}^{i-1}, \mathbf{h}^\ell, \mathbf{h}^\ell) \\
\hat{\mathbf{d}}^i &= \text{FFN}(\hat{\mathbf{D}}) \\
\mathbf{d}^i &= \text{Adapter}(\hat{\mathbf{d}}^i)
\end{aligned}$$

Here,  $\ell$  refers to the final encoder layer, and  $\text{MHA}(\mathbf{Q}, \mathbf{K}, \mathbf{V})$  denotes the standard multi-head attention mechanism (Vaswani, 2017), where  $\mathbf{Q}$ ,  $\mathbf{K}$ , and  $\mathbf{V}$  are the queries, keys, and values, respectively.

## B.2 Implementation Details

The architecture of the SeamlessM4T medium incorporates a speech encoder that has 12 conformer layers, while both the text encoder and text decoder consist of 12 Transformer blocks, with a model dimension of  $D_1 = 1024$ . In our experiments, adapters were introduced after each encoder conformer layer and the decoder Transformer layer. These adapters project the original  $D_1$ -dimensional features into a reduced intermediate space of dimension  $D_2$ , apply a GeLU non-linear activation function (Hendrycks and Gimpel, 2023), and then project the features back to  $D_1$ . The projected layer dimension on the adapters is  $D_2 = 2048$ . The value of  $D_2$  controls the number of trainable parameters, with smaller values of  $D_2$  reducing parameter count. With  $D_2$  set to half of  $D_1$ , this setup introduced 100M trainable parameters while keeping the rest of the model frozen.

All the fine-tuning experiments were conducted using the SeamlessM4T codebase (Communication et al., 2023) released by Meta AI using NVIDIA RTX A6000 GPUs. The experiments were conducted over 20 epochs, utilizing a batch size of 8 and a learning rate of  $5 \times 10^{-6}$ . All the reported results throughout this study are based on a single fixed random seed.

The paraphrase generation using IndicTrans2 and NLLB employs a beam width of 5, while TopK and Nucleus sampling utilize  $K = 50$  and  $P = 0.95$ , respectively.

## C LLM Prompts for Paraphrasing

The paraphrasing prompt given to the Aya model for our very specific paraphrasing task has been stated below:

*Paraphrase the following sentence in **lang**, strictly adhering to these guidelines:*

1. *Maintain the original sentence structure and word order as much as possible.*

2. *Replace at least one word, and aim to replace as many words as feasible with Hindi synonyms or words with similar meanings.*
3. *Do not add extra words or elaborate on the description.*
4. *Preserve named entities (e.g., proper names, places) in their original form.*
5. *Convert ALL numbers to their Hindi word equivalents. This includes dates, years, percentages, and any other numerical values.*
6. *Ensure that all replacements are common Hindi words, avoiding obscure or highly technical terms.*
7. *If a direct Hindi synonym is not available, use a phrase that conveys the same meaning.*
8. *Maintain the original tense and grammatical structure of the sentence.*
9. *If the original sentence contains English words commonly used in Hindi, you may keep them unchanged.*

**IMPORTANT:** Double-check that NO numerical digits remain in your paraphrase. All numbers must be written out in Hindi words.

*Examples: Some Hindi examples with the required paraphrases were provided*

## D Some Qualitative examples

### D.1 Model Outputs

Table 6 depicts examples of phrases that were acceptable for human annotation but would have incurred penalties on the use of other metrics. It can be observed that the model outputs differ from the ground truth due to native spellings of English words, whether compound words are connected or not, and semantically similar but linguistically different words and phrases. Such errors get penalized harshly by metrics like WER.

### D.2 Paraphrases

Table 7 shows examples of sentences and their corresponding paraphrases generated via round-trip translation, where word order has been preserved to ensure semantic alignment. These were used as a guideline to create the paraphrasing prompt of the LLM. We require paraphrases where word order does not change much and where synonyms and semantically similar but linguistically different phrases are used frequently.



Language	ASR	AMPS <sub>T</sub>	Meaning	Explanation
Marathi	aaiskrim	aayskrim	icecream	Different native spelling of english word
	aplya sarkhya	aplyasarkhya	like ours	Compound words joined together
	tyoob	tyub	tube	Different native spelling of english word
Hindi	baaki kuch nahi	aur kuch nahi	nothing else	Semantically similar phrases
	bhajansangraha	bhajan sangraha	prayer collection	Compound words separated
	manobhavon	bhavanaon	sentiments	Semantically similar words

Table 6: Examples of semantically similar and linguistically different phrases and words

Language	Ground Truth	Paraphrase
Marathi	plij mala sagla informashun dya	krupaya tumhi mala sarva mahiti dya
	aani ashya bimarina rokhne	aani ashya roganpasun bachav karne
Hindi	draiving karte samay mobail fon ka yuj nahi kare	gaadi chalte samay mobail fon ka upyog na kare
	kareer banana pasand karunga iska pramukh kaaran	kareer banana chahunga jiska mukhya kaaran

Table 7: Examples demonstrating the ideal paraphrases for AMPS.

## E Paraphrase Evaluation Metrics

1. **Word Error Rate (WER)** measures the number of mistakes in transcription as a ratio of the number of words. These errors could be substitutions, insertions or deletions.

$$\text{WER} = \frac{\text{Substitutions} + \text{Inclusions} + \text{Deletions}}{\text{Words in Reference Text}} \quad (2)$$

2. **METEOR** (Banerjee and Lavie, 2005) is used for evaluating of machine translation quality. It has also previously been used for evaluating paraphrase quality (Shen et al., 2022b). It aligns words in the candidate and reference translations based on word level matches, including same meaning words and stemming.
3. **BERTScore** (Zhang et al., 2020) evaluates the similarity between two texts by using BERT embeddings (Devlin et al., 2019) (Bidirectional Encoder Representations from Transformers). It captures contextual meaning and semantics by computing the cosine similarity between token embeddings from a reference sentence and a candidate sentence. We used AI4Bharat’s IndicBERT (Kakwani et al., 2020) for our BERTScores.
4. **Other metrics** like PARAScore (Shen et al.,

2022b), BBScore (Shen et al., 2022a), LATTEScore (Tomanek et al., 2024) and ROUGE (Patil et al., 2022) have been used in the past for evaluation of paraphrases.

## F AMPS for Read Speech

Table 8 depicts AMPS for Marathi, Malayalam, and Kannada using all the read speech of the IndicVoices (Javed et al., 2024a) dataset. Training sets of Kannada, Malayalam, and Marathi were of duration 2.64, 2.01, and 4.84, respectively. All validation sets were of a half-hour duration. It can be observed that AMPS<sub>T</sub> performs the best for Marathi, Malayalam, and Kannada round-trip translated read speech.

## G 5 hour AMPS for Other languages

Table 9 depicts the two different round-trip translation methods used for AMPS for 5 hours each of mixed Marathi, Malayalam and Kannada speech. It can be observed that the two methods have comparable performance, with normal round-trip translation performing slightly better than the top-K and nucleus (top-P) setting.

## H Details of Human Evaluation

Human evaluation was outsourced to an annotation company based in India, and INR 45 was paid for

Language	Paraphrase Type	Baseline	Read Speech RT Trans		
	Configuration		ASR	AMPS	AMPS <sub>τ</sub>
Marathi	WER ↓	38.65	34.04	32.30	<b>31.25</b>
	METEOR ↑	59.84	67.26	68.83	<b>70.04</b>
	BERTScore ↑	81.01	87.71	88.65	<b>89.18</b>
Malayalam	WER ↓	56.15	55.38	55.17	<b>54.58</b>
	METEOR ↑	43.69	45.85	45.59	<b>46.22</b>
	BERTScore ↑	84.35	85.72	<b>86.01</b>	85.99
Kannada	WER ↓	69.29	61.86	61.3	<b>59.64</b>
	METEOR ↑	31.13	38.95	39.80	<b>40.63</b>
	BERTScore ↑	76.65	82.48	82.52	<b>83.04</b>

Table 8: Comparison of ASR performance for pure ASR, AMPS and AMPS<sub>τ</sub> with round-trip translated (RT Trans) read-speech data for Marathi, Malayalam and Kannada

every audio. Each sentence was given a maximum score of 5 for perfect transcription. In cases of erroneous transcriptions, 0.5 points were deducted for every instance of a minor error, and 1 point was deducted for every instance of a major error. Minor errors included small character errors or tense changes that led to wrong grammar. Major errors included wrong transcriptions, missed words, and wrongly spelled native words. The annotators were instructed to give no penalty for incomprehensible audio, varying native spellings of English words or proper nouns, semantically similar but linguistically different words, and broken or connected compound words.

## I Paraphrase Supervision for Purely Speech-to-Text Models

To provide a comparison for our multimodal model technique, we propose an alternative approach involving pretraining and finetuning for purely speech-to-text ASR models. The hypothesis is that training an ASR model first on speech paired with paraphrased transcripts, followed by finetuning it on speech with original transcripts, will result in a model that is more robust to mispronunciations and noisy inputs. By learning to associate unclear or imprecise utterances with semantically similar phrases, this model should outperform one trained exclusively on ground-truth labels when evaluated on noisy test sets despite exposure to similar amounts of data. To support our hypothesis, we used the Whisper ASR model trained sequentially using paraphrased transcripts followed by the

Language	Paraphrase Type	-	Mixed Speech RT Trans		Mixed Speech TK+Nuc RT Trans	
	Configuration	ASR	AMPS	AMPS <sub>τ</sub>	AMPS	AMPS <sub>τ</sub>
Marathi	WER ↓	24.70	<b>24.44</b>	24.60	<b>24.56</b>	24.75
	METEOR ↑	76.66	76.80	<b>77.11</b>	76.50	<b>76.74</b>
	BERTScore ↑	91.77	91.83	<b>92.01</b>	91.59	<b>91.83</b>
Malayalam	WER ↓	47.90	47.11	<b>46.06</b>	46.41	<b>46.27</b>
	METEOR ↑	55.29	<b>55.86</b>	55.82	56.84	<b>56.92</b>
	BERTScore ↑	89.82	<b>90.18</b>	89.96	<b>90.27</b>	90.25
Kannada	WER ↓	46.77	46.53	<b>46.35</b>	46.24	<b>46.22</b>
	METEOR ↑	53.77	54.49	<b>54.80</b>	54.34	<b>54.47</b>
	BERTScore ↑	87.90	87.78	<b>87.92</b>	87.86	<b>87.99</b>

Table 9: Comparison of ASR performance for pure ASR, AMPS and AMPS<sub>τ</sub> for normal round-trip translated (RT Trans) and top K + Nucleus sampled round-trip translated (TK+Nuc RT Trans) mixed data for Marathi, Malayalam, and Kannada

ground truth, with an ASR training objective.

### I.1 Whisper

Whisper (Radford et al., 2022), developed by OpenAI, utilizes a transformer-based encoder-decoder framework suitable for a range of speech-related tasks. The model comprises an audio encoder that processes raw audio inputs, transforming them into log-mel spectrograms. This input is fed into multiple transformer layers designed to capture long-range dependencies within the audio data. The text decoder, operating autoregressively, generates transcriptions from the processed audio features while integrating task-specific tokens for seamless task-switching among any auxiliary tasks.

### I.2 Experiment and Results

The Whisper model was trained sequentially with 5-hour round-trip translated read speech data in three different ways - training with ground truth training followed by paraphrased training, paraphrase training followed by ground truth training, and finally, ground truth training repeated twice.

The WER (%) values for Hindi read speech were 87.68 for direct inference, 42.33 for ground truth - ground truth training, 47.34 for paraphrase - ground truth training and 43.78 for ground truth - paraphrase training. Since pure ground truth training WER is the best, we chose not to proceed with this experiment as this strongly supports that multi-modality of a model is essential for AMPS.