



Svarah: Evaluating English ASR Systems on Indian Accents

Tahir Javed^{1,2}, Sakshi Joshi², Vignesh Nagarajan², Sai Sundaresan², Janki Nawale², Abhigyan Raman², Kaushal Bhogale^{1,2}, Pratyush Kumar^{1,2,3}, Mitesh M. Khapra^{1,2}

¹Indian Institute of Technology Madras, India

²AI4Bharat, India

³Microsoft, India

{tahirjmakhdoomi, sakshijcom, vignesh.vn.nagarajan, saisundaresan01}@gmail.com, janki@ai4bharat.org, ramanabhigyan@gmail.com, {cs22d006, pratyush, miteshk}@cse.iitm.ac.in

Abstract

India is the second largest English-speaking country in the world with a speaker base of roughly 130 million. Thus, it is imperative that automatic speech recognition (ASR) systems for English should be evaluated on Indian accents. Unfortunately, Indian speakers find a very poor representation in existing English ASR benchmarks such as LibriSpeech, Switchboard, Speech Accent Archive, etc. In this work, we address this gap by creating *Svarah*, a benchmark that contains 9.6 hours of transcribed English audio from 117 speakers across 65 geographic locations throughout India, resulting in a diverse range of accents. *Svarah* comprises both read speech and spontaneous conversational data, covering various domains, such as history, culture, tourism, etc., ensuring a diverse vocabulary. We evaluate 6 open source ASR models and 2 commercial ASR systems on *Svarah* and show that there is clear scope for improvement on Indian accents.

Index Terms: non-native speech recognition, Indian accents, diversity, and inclusion

1. Introduction

Recent advances in Automatic Speech Recognition have demonstrated that current models have achieved human parity in transcription [1, 2]. For example, state-of-the-art systems from Google and Microsoft have reported Word Error Rates (WERs) as low as 4.9%¹ and 5.1% [2] respectively. Similarly, open-source models such as Meta’s Wav2vec2 [3] and OpenAI’s Whisper [1] have reported WERs of 1.8% and 2.7% on the LibriSpeech benchmark. However, these benchmarks only contain data from native speakers and a more robust evaluation using benchmarks with a diverse representation of accents is missing. Indeed, recent works [4] have shown that state-of-the-art ASR models perform significantly worse on non-native speakers than on native speakers when evaluated on datasets with a diverse range of accents, such as the Speech Accent Archive (SAA). However, the SAA benchmark has its limitations as (i) it only contains read speech and (ii) it does not have a good representation of Indian speakers (only 59/1500 speakers).

In this work, we focus on creating a robust benchmark for evaluating ASR systems for Indian accents. India is recognized for its rich linguistic diversity, with 22 official languages, 122 major languages, and 1599 other languages in the country, as per the Census of 2011². Despite the abundance of regional languages, English continues to be a language of importance in India, serving as the official language of the Indian government, the Supreme Court of India, and the primary medium

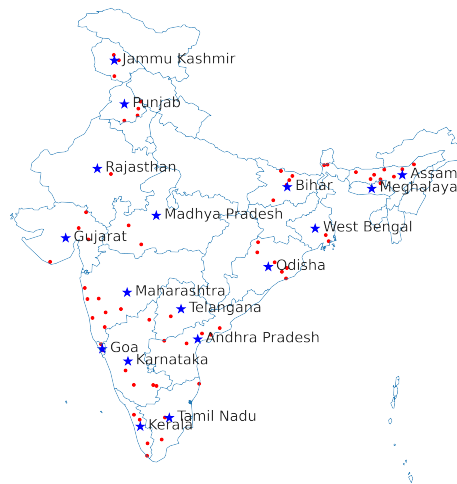


Figure 1: *Svarah* contains data from 117 speakers spanning 65 districts across 19 states in India and speaking 19 out of the 22 constitutionally recognized languages.

of instruction in higher education institutions. Current estimates show that approximately 10% or about 130 million people in India speak English, making it the second largest English-speaking country in the world³. In addition, the accents of English speakers in India vary significantly due to the influence of regional languages. For instance, native speakers of Dravidian languages, such as Tamil and Malayalam, have distinct intonation patterns and pronunciation styles compared to native speakers of Indo-Aryan languages, such as Hindi, Marathi, Gujarati, etc. Even within the Indo-Aryan family, there are significant variations in the accents of speakers from the country’s northern, western, and eastern parts. Hence, a robust English ASR benchmark covering the accent diversity of India is needed.

To address this gap, we introduce *Svarah*, an ASR benchmark containing Indian-accent English data. To build *Svarah*, we recruited speakers from diverse geographical locations spanning the length and breadth of India. Figure 1 shows the distribution of speakers across different parts of the country covering 65 districts across 19 states in India. The benchmark contains a total of 9.6 hours of transcribed data collected from 117 speakers having diverse language backgrounds resulting in different accents. More specifically, the collective set of native languages spoken by the speakers covers 19 of the 22 constitutionally rec-

¹<https://tinyurl.com/49y833xn>

²<https://censusindia.gov.in/nada/index.php/catalog/42561>

³https://en.wikipedia.org/wiki/List_of_countries_by_English-speaking_population

ognized languages of India, belonging to 4 different language families. The dataset includes both read speech and spontaneous conversational data, covering a variety of domains such as history, culture, tourism, government, sports, etc. It also contains data corresponding to popular use cases such as ordering groceries, making digital payments, and using government services (e.g., checking pension claims, checking passport status, etc.). The resulting diversity in vocabulary as well as use cases allows a more robust evaluation of ASR systems for real-world applications.

Using *Svarah*, we evaluate 6 state-of-the-art English ASR models based on different paradigms. These include (i) End-to-end models such as OpenAI’s Whisper [1] which is trained on weakly supervised web-scale data (438K hours) curated from YouTube which is noisy but very diverse (ii) Self-supervised models such as Wav2Vec2, HuBERT, WavLM, and Data2Vec which use a large amount of unlabeled data for pre-training and relatively smaller amounts of clean supervised data (960 hours) (iii) Conformer based model [5] released by NVIDIA which is trained on a large amount of manually labeled data (24.5K hours). We also evaluate two commercial models offered by Google and Microsoft. We observe that the Whisper model which is trained on large amounts of diverse weakly supervised data performs significantly better than the other open-source models mentioned above. Further, two of the six open-source models perform significantly better than both the commercial models and exhibit very little standard deviation across different Indian accents. The benchmark as well as our code has been made publicly available⁴ and will help in a more robust evaluation of English ASR systems on a diverse set of Indian accents.

2. Related Work

There are several datasets available for training and evaluating English ASR systems. A few popular ones which are freely available under permissible licenses include Librispeech [6], Switchboard-1 Dataset [7], WSJ-0 and WSJ-1 [8], VCTK [9], VoxPopuli [10] and Mozilla Common Voice [11]. However, none of these have a good representation of Indian speakers with diverse accents. For example, Librispeech, Switchboard, and the two WSJ datasets have no speakers residing in India. Similarly, VoxPopuli mainly contains European Parliament event recordings and is unlikely to have a representation of a diverse set of Indian speakers. There are a few datasets that specifically address the problem of evaluating ASR systems on multiple accents. These include the Speech Accent Archive (SAA) dataset and The Foreign Accented English (FAE) dataset [12]. SAA only has 59 Indian speakers contributing a total of 26 minutes of data. Similarly, FAE only consists of utterances spoken by native Hindi and Tamil speakers contributing a total of 0.5 hours and 1 hour of data respectively, and does not capture the large linguistic diversity of India. There are three datasets, viz., the IITM and NPTEL datasets⁵ and the AccentDB dataset [13] which specifically focus on Indian accent English. However, the IITM and NPTEL datasets contain data only from the STEM domain and the speakers have good fluency in English given that they come from highly educated backgrounds (professors at premier higher education institutes in India). The AccentDB dataset only contains 8 speakers belonging to 4 Indian languages. In contrast to these benchmarks, *Svarah* contains

more diverse content from a variety of domains and use cases. Further, it is collected from a diverse set of speakers across India speaking 19 different native languages.

3. Svarah

We now describe the various steps involved in creating *Svarah*. **Recruitment of diverse speakers.** Through a network of professional translators, we contacted speakers who were fluent in English but were native speakers of a regional Indian language. For each of the 19 languages, we recruited 3-5 bilingual speakers who spoke English and one of the constitutionally recognized Indian languages resulting in a total of 117 speakers. Of these, 54 were men and 63 were women. We also ensured that we had a roughly equal number of speakers belonging to the following age groups 18-30, 30-45, 45-60, and 60+. The speakers also came from different educational backgrounds (arts, commerce, science) with different levels of education (graduates, post-graduates, PhDs). The task was clearly explained to the speakers, and they were informed that the data was being collected to build and evaluate speech models. Their voice samples were recorded only if the speakers willingly agreed to participate in the task and signed a consent form to this effect.

Collection of multi-domain text data. A part of the data included read speech which required participants to read a piece of text shown to them. To ensure that the text being read covered diverse vocabulary, we collected text from Wikipedia belonging to multiple domains. These domains were identified using the “Category” information available for Wikipedia articles. In total, we collected 1k sentences from 9 domains, viz. health, entertainment, culture, geography, history, business, news, sports, and tourism. Each participant was asked to read 5 sentences randomly chosen from this collection while ensuring that no two participants got the same sentence.

Creation of prompts for collecting extempore data. The participants were asked to fill out a form, where, in addition to meta-data such as age, district of residence, native language, etc. they were also asked to select (i) topics of interest and (ii) specific domains about which they could talk. We considered 28 topics of interest such as painting, cooking, gardening, knitting and stitching, traveling, etc., and the same 9 domains listed above. For each topic of interest, we created a few simple questions that any participant could answer, such as, “What inspired you to take up drawing?”, “What are the dishes you like to cook and tell us the recipe of your favorite dish?”. Similarly for each domain, we created simple questions which anyone interested in that domain could answer. For example, someone interested in Entertainment should be able to answer the following question “What is your favorite movie or TV serial and why do you like it?”. Each participant was shown four such questions for each of the selected topics of interest and domains and was required to answer these questions using extempore speech.

Creation of prompts for different use cases. One of the goals was to create a benchmark that also contains utterances that one would encounter in popular everyday use cases. We considered three such use cases, viz., ordering grocery items, making digital payments, and using government services. For each of these use cases, we created prompts such as, “Add 2kg Tomato, 1 Cadbury Chocolate to my shopping cart”, “Pay 1000 rupees from my SBI bank number 123456789 to UPI ID 9876543210@paytm” and so on. Each participant was shown five such prompts for each of these use cases.

Creation of a mobile app for data collection. Given the diverse locations from which we intended to collect data, we

⁴<https://github.com/AI4Bharat/Svarah.git>

⁵<https://sites.google.com/view/englishasrchallenge/home>

needed a tool that allowed us to distribute and monitor work remotely. To do so, we used Microsoft’s open-source crowd-sourcing platform called Karya which is available as an android application and has already been used by other teams in the past to collect voice data [14, 15]. Once a participant fills out the participation form, they can install Karya and log in with their mobile number. On logging in, the participant will see different tasks corresponding to (i) read speech (ii) questions on topics of interest (iii) questions on domains of interest (iv) prompts on everyday use cases described above. Once they enter a task, they can read the prompt and press the record button to start recording their response to the prompt. Once done, the participant can press the “stop” button, at which point the app automatically replays the audio recorded by the participant. The participants were instructed to listen to the audio to ensure that (i) there was no background noise and (ii) the recording was clearly audible even at 50% of the volume of the device.

Transcription of voice samples. To transcribe the data we recruited undergraduate and postgraduate students who work on various research projects in our institute and are fluent in English (having done their schooling as well a higher education in English). The transcribers were asked to verify the recorded samples to check for audio quality and to also ensure that the responses were on topic (*i.e.*, the participants were responding correctly to the prompts). The transcribers then segmented the data into logical segments (typically, sentences) and then transcribed these segments using a modified instance of Label Studio. Given the extempore and conversational nature of the data, they were asked to use the same transcription guidelines as used for the Switchboard corpus.

In summary, *Svarah*⁶ contains a total of 9.6 hours of transcribed data from a total of 117 Indian speakers from 19 different languages covering 65 districts in 17 Indian states. It contains 1.4, 6.4, and 1.7 hours of (i) read speech, (ii) extempore speech on topics of interest, and (iii) utterances from everyday use cases, respectively. The share of extempore speech is larger as participants tend to give longer responses to questions about their topics of interest as opposed to everyday use cases where the responses are to the point.

4. Methodology

We evaluate a total of 8 models on *Svarah*. Since *Svarah* only contains data from non-native (L2) speakers, to show the contrast with the performance of native (L1) speakers, we consider the LibriSpeech dataset and the Speech Accent Archive (SAA) dataset. In particular, we consider the data from L1 speakers in SAA and evaluate the models on this subset of the data (SAA-L1). Of the 8 models we considered, 6 models are open-source and are publicly available, whereas 2 models are commercially available, one each from Google and Microsoft. From the open-source category, we chose models based on different paradigms as listed below. We do not use an external language model with any of the models listed below as the idea was to evaluate the ability of the acoustic model to handle diverse accents.

End-to-End models. Here, we consider OpenAI’s Whisper model which has been trained on 680K hours of weekly supervised multilingual data of which 438K hours of data is in English. One advantage of this model is the huge diversity in the training data which perhaps makes it better suited for handling diverse accents. The downside, of course, is that the training data is not entirely gold-standard but weakly supervised. We

⁶*Svarah* means accent in Sanskrit

	# Params.	Svarah	SAA_L1	Libri
Whisper _{base}	74M	13.6	2.9	4.2
Whisper _{medium}	769M	8.3	1.7	3.1
Whisper _{large}	1550M	7.2	1.6	2.7
Wav2Vec2 _{large}	317M	24.9	3.1	1.8
HuBERT _{large}	316M	25.6	3.2	2.0
WavLM _{large}	300M	33.7	9.2	3.4
Data2Vec _{large}	313M	24.5	2.5	1.8
Conformer _{large}	120M	14.6	1.1	2.1
Azure _{US}	-	20.9	24.2	-
Azure _{IN}	-	21.3	30.1	-
Google _{US}	-	30.0	16.8	-
Google _{IN}	-	20.7	63.7	-

Table 1: WERs of different models on (i) *Svarah* that contains data from Indian speakers & (ii) SAA_L1, LibriSpeech Clean (Libri) which contain data from native English speakers.

consider the base, medium and large models available as a part of OpenAI’s official release on HuggingFace.

Self-supervised models. Here, we consider models such as Wav2Vec2, HuBERT, WavLM, and Data2Vec which are pre-trained on large amounts of unsupervised data and then finetuned on a relatively smaller amount of supervised data. We used the publicly available versions of these models from HuggingFace, which were pretrained on 60k hours of LibriVox data and then finetuned on 960h of Librispeech data except WavLM where the authors used a 100-hour clean subset of Librispeech for finetuning the model.

Conformer based models. All the models described above are transformer-based models. We also consider the conformer-based model trained with CTC loss as released by NVIDIA. This model does not use any unsupervised pre-training but is trained on a larger amount of gold standard data collated from 11 different sources, resulting in a total of 24,500 hours of training data. We hoped that given the diverse and relatively cleaner and larger training data, this model might perform well on diverse accents.

From the commercially available models, we tested Google⁷ and Azure⁸ using their provided SDKs. In addition to evaluating the international version (en-US) of their ASR models, we also evaluate their India-specific models (en-IN).

5. Results & Discussions

The results of our experiments are summarised in Tables 1 and 2. Below, we discuss the main observations from these results.

Poor performance on non-native speakers. Referring to Table 1, we observe that for all the 6 open source models there is a significant gap between the performance of the models on *Svarah* and L1 data (SAA_L1 and LibriSpeech). More specifically, the gap between the WERs on *Svarah* and SAA_L1 ranges from 5.65 - 43.03. Similarly, the gap between the WERs on *Svarah* and LibriSpeech ranges from 4.5 - 30.3, clearly indicating that the claims of human parity need to be revisited by evaluating on more diverse benchmarks like *Svarah*. Note that the medium and large variants of Whisper give single-digit

⁷<https://cloud.google.com/speech-to-text>

⁸<https://azure.microsoft.com/en-us/products/cognitive-services/speech-to-text>

	W_{large}	W2V2	HuB	D2V	C_{large}	G_{IN}	AZUS
Split by accents							
ne	9.8	34.4	35.2	32.6	20.5	27.5	33.1
brx	11.6	29.9	30.5	29.0	18.5	26.7	38.3
as	10.1	32.3	33.8	31.6	21.4	20.6	18.1
doi	8.0	29.1	28.5	26.8	16.5	21.1	21.6
sd	7.3	27.2	26.5	26.5	16.8	19.1	21.0
ml	8.1	23.7	26.5	23.0	14.7	24.9	20.4
kn	6.6	27.2	27.1	27.0	14.6	17.7	14.2
mr	6.8	24.6	26.0	26.2	15.0	18.6	15.5
bn	7.6	24.1	24.8	25.1	14.1	18.5	16.6
pa	6.7	22.3	22.3	22.6	13.3	17.4	25.2
ur	6.2	22.3	23.3	21.8	14.2	21.5	19.7
te	7.5	23.4	23.1	22.6	14.3	18.8	17.6
ks	6.9	20.3	21.2	20.3	12.6	27.8	18.3
kok	6.4	20.1	21.3	19.7	11.9	19.2	25.3
or	6.2	24.5	24.9	24.8	12.1	17.5	13.4
guj	6.0	23.3	24.0	22.6	12.7	13.1	17.8
hi	5.3	20.7	20.8	20.4	10.9	19.7	21.0
mai	4.5	20.7	21.9	20.7	11.6	15.4	19.6
ta	5.3	17.8	18.8	18.6	10.7	17.4	9.9
Split by type of data							
Re	6.2	22.8	23.3	21.7	12.6	21.3	24.0
Ex	7.4	23.4	25.3	28.0	14.0	18.4	10.1
UC	11.2	34.6	35.4	33.7	23.3	19.8	15.7

Table 2: The WER of different models on different splits of Svarah based on (i) accents and (ii) types of data (Re: read speech, Ex: extempore speech, UC: use cases). Model names are abbreviated as W_{large} : Whisper; W2V2: Wav2Vec2, HuB: HuBERT, D2V: Data2Vec, G_{IN} : Google_{IN}, AZUS: Azure_{US}.

WERs which is encouraging but, given the large size of these models, it may be difficult to deploy them in the Indian context. Surprisingly, both the Azure models perform better on Svarah than on the L1 data. It is hard to comment on why this is the case as the details of Azure’s models are not known.

Comparison of different models. Again referring to Table 1, we observe that the Whisper family of models gives the best performance on Svarah with the smallest gap in WERs on Svarah and L1 data. We hypothesize that this is due to the large diversity in the weakly supervised training data used by these models which makes them more robust to diverse accents. In other words, the trade-off between the quality and volume of the training data works favorably. The next best-performing model is the NVIDIA Conformer-based model which was again trained on a large amount of gold standard data collated from different sources (again reinforcing the importance of diversity in training data). The self-supervised models (Wav2Vec2, HuBERT, WavLM and Data2Vec) perform poorly, indicating that pre-training on large amounts of diverse unlabelled data is not as useful as finetuning on diverse data. WavLM which is trained on the smallest amount of training data (100 hours) gives the worst performance. Lastly, the two commercial models are outperformed by the open-source Whisper and NVIDIA Conformer models, making a clear case for using open-source models.

Performance across different accents. Table 2 shows the performance of the models split by the native language spoken by the speakers. Due to space constraints, we consider only the top-7 models from Table 1. We observe that the best-performing




Figure 2: Correlation between the WERs of different models across different accents.

models, viz., Whisper-Large and NVIDIA Conformer, have very low standard deviation across different accents (1.8 and 3.0 respectively). In general, the models perform poorly for speakers who are native speakers of Assamese, Bodo, and Nepali which happen to be low-resource languages from the North-Eastern region of India. Figure 2 shows that there is a high correlation between the performance of the open-source models across accents, indicating that in general, the relative performance of models across different accents is similar. Further, the correlation between the 3 self-supervised models is higher given the similar pre-training and fine-tuning data used by these models. Both the commercial models do not correlate well with any of the open source models perhaps due to different training data (including proprietary data) and design choices.

Performance across different types of data. As mentioned earlier, Svarah contains read speech, extempore data and utterances corresponding to everyday use-cases. Referring to the last section of Table 2, we observe that except for the two commercial models, all models perform better on read and extempore data as compared to the data corresponding to use cases. This is mainly because the data corresponding to everyday use cases contains many entities such as brand names, bank names, food items, document IDs, and so on. Recognizing such entities when spoken in Indian accents is hard for existing ASR systems. Even the Whisper_{large} model performs poorly on the utterances from everyday use cases with a WER of 11.2, indicating that there is much scope for improvement.

6. Conclusion

In this work, we addressed the lack of Indian-accented English data in existing speech recognition benchmarks by introducing Svarah, a dataset containing 9.6 hours of transcribed data collected from 117 speakers across India, covering a broad range of accents and linguistic backgrounds. Our dataset includes both read speech and spontaneous conversational data, covering a variety of domains such as history, culture, tourism, government, sports, etc. We evaluated six open-source and two commercial ASR models on this dataset and found that the recently released Whisper model outperformed commercial models by a significant margin. Our dataset and all evaluation scripts have been made publicly available. The creation of Svarah will enable a more robust evaluation of ASR systems for real-world applications in India, where English is an important language, and a diverse range of accents exists due to the influence of regional languages.

7. References

- [1] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," *CoRR*, vol. abs/2212.04356, 2022. [Online]. Available: <https://doi.org/10.48550/arXiv.2212.04356>
- [2] W. Xiong, J. Droppo, X. Huang, F. Seide, M. L. Seltzer, A. Stolcke, D. Yu, and G. Zweig, "Toward human parity in conversational speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 12, pp. 2410–2423, 2017.
- [3] A. Baeovski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., 2020. [Online]. Available: <https://proceedings.neurips.cc/paper/2020/hash/92d1e1eb1cd6f9fba3227870bb6d7f07-Abstract.html>
- [4] S. Hollands, D. Blackburn, and H. Christensen, "Evaluating the performance of state-of-the-art ASR systems on non-native english using corpora with extensive language background variation," in *Interspeech 2022, 23rd Annual Conference of the International Speech Communication Association, Incheon, Korea, 18-22 September 2022*, H. Ko and J. H. L. Hansen, Eds. ISCA, 2022, pp. 3958–3962. [Online]. Available: <https://doi.org/10.21437/Interspeech.2022-10433>
- [5] A. Gulati, J. Qin, C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented transformer for speech recognition," in *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*, H. Meng, B. Xu, and T. F. Zheng, Eds. ISCA, 2020, pp. 5036–5040. [Online]. Available: <https://doi.org/10.21437/Interspeech.2020-3015>
- [6] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An asr corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5206–5210.
- [7] J. J. Godfrey, E. Holliman, and J. McDaniel, "Switchboard: telephone speech corpus for research and development," [*Proceedings*] *ICASSP-92: 1992 IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 1, pp. 517–520 vol.1, 1992.
- [8] L. Drude, J. Heitkaemper, C. Boeddeker, and R. Haeb-Umbach, "Sms-wsj: Database, performance measures, and baseline recipe for multi-channel source separation and recognition," *ArXiv*, vol. abs/1910.13934, 2019.
- [9] J. Yamagishi, C. Veaux, and K. MacDonald, "Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit (version 0.92)," 2019.
- [10] 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 *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Online: Association for Computational Linguistics, Aug. 2021, pp. 993–1003. [Online]. Available: <https://aclanthology.org/2021.acl-long.80>
- [11] R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, "Common voice: A massively-multilingual speech corpus," in *International Conference on Language Resources and Evaluation*, 2019.
- [12] T. Lander, "CSLU: Foreign Accented English Release 1.2," 2022. [Online]. Available: <https://doi.org/10.5683/SP2/K7EQTE>
- [13] A. Ahamad, A. Anand, and P. Bhargava, "Accentdb: A database of non-native english accents to assist neural speech recognition," in *Proceedings of The 12th Language Resources and Evaluation Conference*. Marseille, France: European Language Resources Association, May 2020, pp. 5353–5360. [Online]. Available: <https://www.aclweb.org/anthology/2020.lrec-1.659>
- [14] B. Abraham, D. Goel, D. Siddarth, K. Bali, M. Chopra, M. Choudhury, P. Joshi, P. Jyothi, S. Sitaram, and V. Seshadri, "Crowdsourcing speech data for low-resource languages from low-income workers," in *Proceedings of The 12th Language Resources and Evaluation Conference, LREC 2020, Marseille, France, May 11-16, 2020*, N. Calzolari, F. Béchet, P. Blache, K. Choukri, C. Cieri, T. Declerck, S. Goggi, H. Isahara, B. Maegaard, J. Mariani, H. Mazo, A. Moreno, J. Odijk, and S. Piperidis, Eds. European Language Resources Association, 2020, pp. 2819–2826. [Online]. Available: <https://aclanthology.org/2020.lrec-1.343/>
- [15] T. Javed, K. S. Bhogale, A. Raman, A. Kunchukuttan, P. Kumar, and M. M. Khapra, "Indicsuperb: A speech processing universal performance benchmark for indian languages," *CoRR*, vol. abs/2208.11761, 2022. [Online]. Available: <https://doi.org/10.48550/arXiv.2208.11761>