



HAL
open science

Speech Foundation Models on Intelligibility Prediction for Hearing-Impaired Listeners

Santiago Cuervo, Ricard Marxer

► **To cite this version:**

Santiago Cuervo, Ricard Marxer. Speech Foundation Models on Intelligibility Prediction for Hearing-Impaired Listeners. ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Apr 2024, Seoul, France. pp.1421-1425, 10.1109/ICASSP48485.2024.10447907 . hal-04592508

HAL Id: hal-04592508

<https://hal.science/hal-04592508v1>

Submitted on 4 Sep 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

SPEECH FOUNDATION MODELS ON INTELLIGIBILITY PREDICTION FOR HEARING-IMPAIRED LISTENERS

Santiago Cuervo, Ricard Marxer

Université de Toulon, Aix Marseille Univ, CNRS, LIS, France

ABSTRACT

Speech foundation models (SFMs) have been benchmarked on many speech processing tasks, often achieving state-of-the-art performance with minimal adaptation. However, the SFM paradigm has been significantly less explored for applications of interest to the speech perception community. In this paper we present a systematic evaluation of 10 SFMs on one such application: Speech intelligibility prediction. We focus on the non-intrusive setup of the *Clarity Prediction Challenge 2* (CPC2), where the task is to predict the percentage of words correctly perceived by hearing-impaired listeners from speech-in-noise recordings. We propose a simple method that learns a lightweight specialized prediction head on top of frozen SFMs to approach the problem. Our results reveal statistically significant differences in performance across SFMs. Our method resulted in the winning submission in the CPC2, demonstrating its promise for speech perception applications.

Index Terms— Foundation models, speech perception, intelligibility prediction, hearing aids.

1. INTRODUCTION

Foundation models (FMs) are deep neural networks trained on large diverse datasets that can be applied to a wide range of downstream tasks with little or no adaptation. They have led to great progress in natural language processing problems on both text and speech [1, 2, 3, 4, 5], often reaching state-of-the-art performance with little task-specific training.

In the speech domain, which is the focus of our work, several studies have benchmarked speech FMs (SFMs) on their capacity to extract phonetic, semantic, speaker-related and paralinguistic information [6, 7]. However, there is a lack of systematic benchmarks on speech perception tasks, in which the objective is not to decode the content of the speech signal, but to model how the signal will be decoded by a listener.

We believe that SFMs have significant potential to drive progress on speech perception. Our belief is based on two

facts: *a)* Representations learned by SFMs can capture semantically meaningful information from both speech [6] and the acoustic background in which it occurs [8], and *b)* Language processing in SFMs and brains shows significant similarities [9]. Based on *a)* we deduce that SFMs can extract the listener-independent information relevant for perception from the whole soundscape. From *b)*, we speculate that such similarities translate in predictive power over perceptual phenomena. We presented evidence for this hypothesis in [10].

In this work we take a step towards benchmarking SFMs for speech perception. We systematically evaluate 10 different SFMs on the *Clarity Prediction Challenge 2* (CPC2) [11], in which the objective is to build machine learning models that take speech-in-noise binaural audio recordings and the audiogram from a hearing-impaired listener as input, and predict the percentage of words that the listener correctly recognizes. Our main contributions are:

- We present a prediction model for the CPC2 leveraging SFMs. The model uses a specialized prediction head that extracts multi-layer features from SFMs, and uses cross-attention to exploit binaural cues. This method resulted in the winning submission for the CPC2.
- We benchmark 10 SFMs on the CPC2 using our specialized predictive model, finding statistically significant performance differences. Additionally, we study SFM ensembles obtained by averaging predictions. We demonstrate that SFMs learn complementary information, as evidenced by over 60% of their ensembles outperforming the best individual model, and all ensembles outperforming all but the single best model.

2. INTELLIGIBILITY PREDICTION MODEL

Our model for intelligibility prediction is depicted in Fig. 1. It consists of a pre-trained frozen backbone that extracts representations from raw binaural speech data, and a head designed

The authors thank the French National Research Agency for their support through the ANR-20-CE23-0012-01 (MIM) grant, and the Institute of Convergence ILCB, supported by grants from France 2030 (ANR-16-CONV-0002) and the Excellence Initiative of Aix-Marseille University (A*MIDEX).

© 2024 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

to predict intelligibility conditioned on the representations obtained from the backbone. We extract features from all layers of the backbone following [6, 5, 8].

2.1. Backbone

For the backbone we consider multiple state-of-the-art SFMs:

Wav2vec 2.0 consists of two components: An encoder and a predictor. The encoder is a strided convolutional network that produces embeddings from the speech input. The predictor is an encoder-only transformer [12] trained contrastively to distinguish positive and negative samples from discretized and masked segments of the encoder’s output. We use four variants: LARGE pretrained wav2vec 2.0 [2] trained on 60k hours of the LibriLight dataset (LL-60k), LARGE pretrained wav2vec 2.0 robust [13] trained on 3k hours of additional diverse data, and their variants fine-tuned (FT) for ASR on 960h of the LibriSpeech dataset (LS-960h).

HuBERT uses an encoder-predictor architecture as in wav2vec 2.0. Instead of a contrastive loss, it is trained with a masked language modeling task similar to BERT [1], but masks continuous segments of the encoder’s output instead of tokens. The targets for the masked prediction are obtained through unsupervised clustering of raw speech features or learned features from earlier iterations of the model. We use the LARGE and X-LARGE variants [3] pre-trained on LL-60K, and their versions fine-tuned for ASR on LS-960h.

WavLM expands on HuBERT with an auxiliary denoising task. Some inputs are corrupted by mixing them with noise or other speech. The model is trained to jointly predict the targets for the masked segments and original uncorrupted speech. This joint pre-training helps the model learn useful representations for a wide range of speech tasks. We use the LARGE variant [5] pre-trained on 94k hours of speech.

Whisper is trained through large-scale multi-task supervised learning on over 680k hours of diverse audio, unlike previous models which were trained using self-supervised learning. The model is an encoder-decoder transformer trained on many speech processing tasks, including multilingual speech recognition, speech translation, spoken language identification, and voice activity detection. Whisper showed robustness to unseen domains, being often competitive with prior fully supervised results but in a zero-shot setting. We use the encoder of the LARGE-v2 variant [4].

2.2. Prediction head

We adapt the transformer across time and layers from [8] for intelligibility prediction by adding an embedding layer to insert listener information, and a cross-attention block between binaural channels to account for binaural effects on intelligibility. We describe below its components.

Dimensionality reduction: The output of the backbone is an $L \times t \times d$ tensor, L being the number of layers of the

SFM, t the length in time of the speech representations sequence, and d the dimension of the representations. Considering the quadratic computational complexity of transformers on the sequence length, we downsample in time by a factor of 20 through average pooling, following [8]. We then project the embeddings from their original d -dimensional space to a 384-dimensional space through a learnable linear projection.

Temporal attention pooling: We use a transformer with bidirectional attention to compress the downsampled sequence along the time axis. As in [1], we append a special CLS embedding at the beginning of the sequence, which after being processed by the transformer can be used as a compressed representation of the whole sequence. Temporal pooling gives an $L \times 384$ tensor.

Audiogram embedding: We use audiograms as listener’s information. Audiograms are a vector in \mathbb{R}^{+8} representing hearing capacity at 8 different frequencies. We project it to \mathbb{R}^{384} through a learnable linear projection and append it along the layer axis, yielding an $(L + 1) \times 384$ -dimensional tensor.

Attention pooling across layers: Next, we apply attention pooling along the layer axis using the same method as in temporal attention. We use the special CLS embedding to summarize the sequence of layer and audiogram embeddings, resulting in a single 384-dimensional representation.

Binaural cross-attention: Each binaural channel is processed by the pipeline described above. To account for non-linear binaural effects in intelligibility [14] we use cross-attention in the attention pooling transformers (see Fig. 1).

Prediction layer: The whole process yields two 384-dimensional representations, one per binaural channel. Both representations are averaged and linearly projected to \mathbb{R} . We apply a sigmoid layer to the output and multiply by 100 to obtain predictions in the range of the target.

3. EXPERIMENTS AND RESULTS

3.1. Experimental setup

3.1.1. Data, metrics, and baselines

CPC2 provides binaural speech signals processed by hearing aids, audiograms from hearing impaired listeners to which the signals are presented, and their responses. The data is split into three *train/test* partitions. The challenge requires training three prediction systems, one per train set. There are 24,630 train and 897 test samples. For each partition, we sample a *dev* set using complex scenes from another partition. Signals are at 32 kHz, we resampled to 16 kHz for the SFMs.

We report average *test* RMSE, the metric used in the CPC2. Significance is tested with a paired Wilcoxon signed-rank test [15]. Baselines are logistic regression on HASPI [16] (CPC2 baseline) and CPC2 SFM-based submissions.

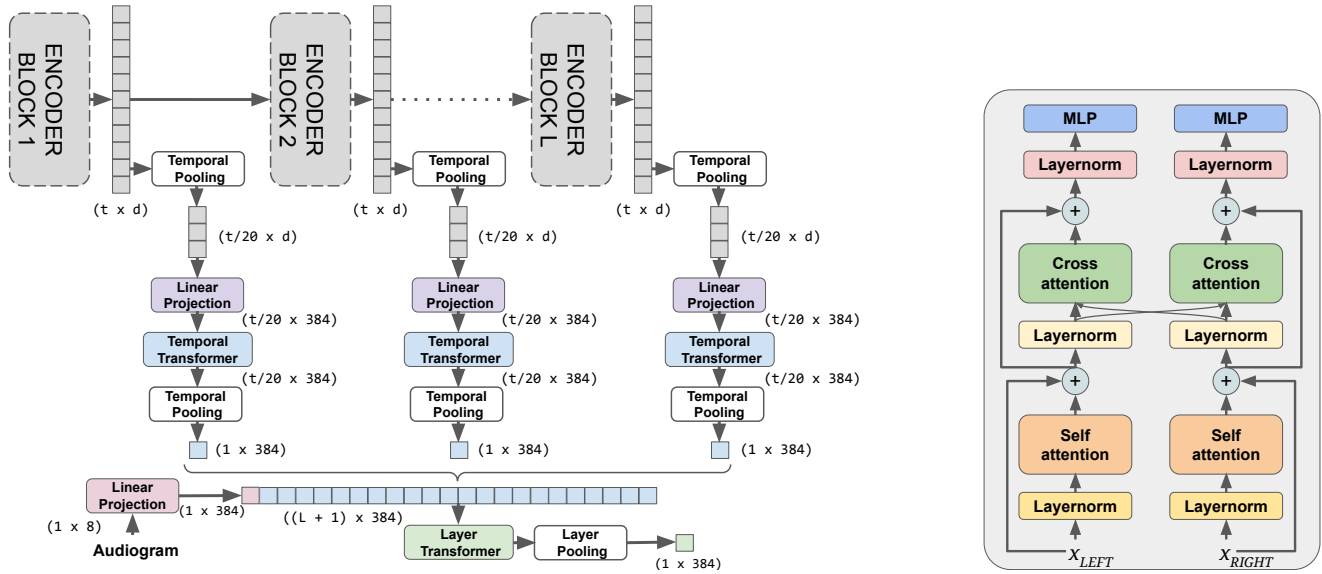


Fig. 1: Intelligibility prediction model architecture. With the exception of frozen backbone blocks (grey), blocks with the same color indicate shared parameters. *Left:* Pipeline applied to each channel of the binaural signal. *Right:* Binaural block used in the temporal and layer transformers. The cross-attention layer enables modeling of non-linear binaural interactions.

3.1.2. Training

All models were trained to minimize a Huber loss for 60k steps using Adam [17] with a learning rate of $3e-5$, $\beta_1 = 0.9$, $\beta_2 = 0.98$, and a batch size of 160. We used a cosine learning rate schedule with a linear warm-up of 2000 steps. To all transformer layers we apply dropout [18] with $p = 0.1$.

3.1.3. Software and computational cost

We use publicly available checkpoints for the backbones. A training run takes close to 9 hours on a single GPU NVIDIA A100-80 GB and it requires about 18.4 GB of GPU memory when using HuBERT X-LARGE or Whisper, and 14.4 GB with other models. We will release our source code at <https://github.com/tiagoCuervo/sfms-cpc2>.

3.2. Results

3.2.1. Backbone performance

Table 1 shows results for different SFM backbones. All backbones outperform the CPC2 baseline. Wav2vec 2.0 FT has the lowest average RMSE, but HuBERT L has the overall lowest. Statistical tests (Table 2) show wav2vec 2.0 FT does not significantly outperform WavLM, Hubert L FT, and Hubert L. These four can be considered the best performing models.

<https://github.com/facebookresearch/fairseq/>,
<https://github.com/microsoft/unilm>, and <https://github.com/openai/whisper>

Table 1: Prediction test RMSE on CPC2 for different backbones using our model. *Top:* across 15 independent runs and of best models chosen by cross validation. *Bottom:* for our submission and baselines.

Backbone	test RMSE ↓		Best on dev
	Min / Mean / Max		
wav2vec 2.0	27.15 / 28.64 / 30.15		28.27
wav2vec 2.0 FT	26.65 / 27.76 / 28.80		26.74
wav2vec 2.0 robust	27.21 / 29.21 / 30.31		28.74
wav2vec 2.0 robust FT	26.75 / 30.05 / 32.39		27.36
HuBERT L	25.05 / 27.89 / 29.52		25.05
HuBERT L FT	26.24 / 27.94 / 30.09		26.92
HuBERT XL	26.69 / 29.47 / 32.80		26.85
HuBERT XL FT	27.40 / 30.61 / 33.50		28.50
WavLM	25.28 / 27.88 / 29.03		25.28
Whisper	26.23 / 28.85 / 30.73		27.83
CPC2 baseline (HASPI)	28.70 / - / -		-
CPC2 E023 (Whisper)	26.34 / - / -		-
CPC2 E002 (Whisper)	25.30 / - / -		-
CPC2 E011 (ours)	25.10 / - / -		25.10

3.2.2. Binaural cross-attention ablation

We evaluated the influence of binaural cross-attention in our predictive model. We conducted ablations with two backbones with significant statistical differences: WavLM and Whisper. The results are presented in Table 3, and show that binaural cross-attention significantly improves performance.

Table 2: Model ranking according to statistical significance with respect to wav2vec 2.0 FT. P -values for a Wilcoxon paired test on the hypothesis of model _{i} having greater RMSE.

model _{i}	Wilcoxon p -val
	RMSE _{i} > RMSE _{wav2vec 2.0 FT}
WavLM	0.598022
HuBERT L FT	0.319336
HuBERT L	0.151398
HuBERT XL	0.047688
wav2vec 2.0	0.037302
wav2vec 2.0 robust	0.037302
Whisper	0.036499
wav2vec 2.0 robust FT	0.027679
HuBERT XL FT	0.006226

Table 3: Effect of binaural cross-attention on performance measured across 15 independent runs.

Model	test RMSE ↓		Wilcoxon p -val RMSE _{with} < RMSE _{w/out}
	with Binaural attention	w./out Binaural attention	
WavLM	27.88 ± 2.28	28.65 ± 1.94	0.021322
Whisper	28.85 ± 2.64	29.46 ± 2.07	0.018101

3.2.3. Our submission to the CPC2

Our submission was an ensemble between the models with the lowest dev RMSE using WavLM and Whisper as backbones. The results we obtained are shown in Table 1. Although our submission was an ensemble, the results show that predictions of single WavLM or HuBERT L backbones would have outperformed others.

3.2.4. Ensemble performance

From our submission we learned that using ensembles of predictions from different backbones improved performance. Therefore, we experimented using ensembles obtained by averaging the predictions from the best performing models for each backbone, chosen by cross validation. Results for the two best and worst ensembles are presented in Table 4, showing that a relative improvement of 5% could be obtained using the best ensemble. It also should be noted that even the worst ensemble outperforms all but the best single model. Furthermore, the whole set of results showed that over 60% of ensembles outperform the best single model.

4. RELATED WORK

Our predictive model was first presented in our submission to the CPC2. With respect to that work, we significantly expanded the amount of experiments, including statistically validated benchmarking of 8 additional SFMs, ablations of binaural cross-attention, and evaluation of SFMs ensembles.

Table 4: Two best (top) and worst (bottom) performing ensembles.

Rank	Ensemble	RMSE ↓
1	HuBERT L+wav2vec 2.0 robust FT	23.86
2	HuBERT L+WavLM	23.88
44	wav2vec 2.0 robust+Hubert XL	25.24
45	wav2vec 2.0 robust+wav2vec 2.0 robust FT	25.25

Moreover, the results here presented outperform our submission and refute the main hypothesis on which we based it (see section 5). Our paper is also related to previous works benchmarking SFMs on the CPC1 [19, 20], but goes beyond them by considering a larger set of SFMs and showing state-of-the-art performance in the most recent edition of the challenge.

5. DISCUSSION

Our results showed that four SFMs performed significantly better than others. All but one of them (wav2vec 2.0 FT) use masked language modeling training, so this could be a factor contributing to their performance. Three out of four of the best models were trained on clean speech only. This refutes the hypothesis behind our CPC2 submission, namely that backbones trained on speech with diverse backgrounds should perform best, and which led us to choose WavLM and Whisper for our submission. Contrary to the trend in other applications, the largest models (Hubert XL and Whisper) performed worse. This could be due to overfitting to extra variance captured by models with higher capacity.

Our experiments with ensembles demonstrated that different SFMs can compensate for each other biases. This suggests that a SFM better suited for intelligibility prediction could be achieved by combining properties from the models studied.

Regarding the design choices for our prediction head, more experiments could be needed to justify them. Binaural cross-attention was shown to have significant impact in performance, however further experiments are required to determine if this is due to exploiting binaural cues, or to gains in expressive power from the extra parameters in the cross-attention block. The effects of downsampling should also be studied, as it could obfuscate temporal fine-structure useful for intelligibility prediction. The role of the audiogram in the models’ predictions remains an interesting question to be studied in future work.

6. CONCLUSIONS

We presented a prediction model leveraging SFMs that obtains state-of-the-art performance in the CPC2. Using that predictive model we benchmarked 10 different SFMs, revealing statistically significant differences in performance. Our experiments also showed that ensemble methods show mean-

ingful improvements, and suggest that different SFMs learn different information useful for intelligibility prediction on hearing-impaired listeners.

7. REFERENCES

- [1] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proc. of NAACL-HLT*, 2019, pp. 4171–4186.
- [2] A. Baevski, 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*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., 2020, vol. 33, pp. 12449–12460.
- [3] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, “HuBERT: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Trans. Audio Speech Lang.*, vol. 29, pp. 3451–3460, 2021.
- [4] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. Mcleavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” in *Proc. of the 40th Int. Conf. on Machine Learning*, A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, Eds. 23–29 Jul 2023, vol. 202 of *Proceedings of Machine Learning Research*, pp. 28492–28518, PMLR.
- [5] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, J. Wu, M. Zeng, X. Yu, and F. Wei, “WavLM: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE J. Sel. Top. in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [6] S. wen Yang, P.-H. Chi, Y.-S. Chuang, C.-I. J. Lai, K. Lakhotia, Y. Y. Lin, A. T. Liu, J. Shi, X. Chang, G.-T. Lin, T.-H. Huang, W.-C. Tseng, K. tik Lee, D.-R. Liu, Z. Huang, S. Dong, S.-W. Li, S. Watanabe, A. Mohamed, and H. yi Lee, “SUPERB: Speech processing Universal PERFORMANCE Benchmark,” 2021.
- [7] T. A. Nguyen, M. de Seyssel, P. Rozé, M. Rivière, E. Kharitonov, A. Baevski, E. Dunbar, and E. Dupoux, “The Zero Resource Speech Benchmark 2021: Metrics and baselines for unsupervised spoken language modeling,” 2020.
- [8] Y. Gong, S. Khurana, L. Karlinsky, and J. Glass, “Whisper-AT: Noise-Robust Automatic Speech Recognizers are Also Strong General Audio Event Taggers,” in *Proc. Interspeech*, 2023, pp. 2798–2802.
- [9] J. Millet, C. Caucheteux, P. Orhan, Y. Boubenec, A. Gramfort, E. Dunbar, C. Pallier, and J.-R. King, “Toward a realistic model of speech processing in the brain with self-supervised learning,” in *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., 2022, vol. 35, pp. 33428–33443.
- [10] S. Cuervo and R. Marxer, “On the Benefits of Self-supervised Learned Speech Representations for Predicting Human Phonetic Misperceptions,” in *Proc. Interspeech*, 2023, pp. 1788–1792.
- [11] S. Graetzer, M. Akeroyd, J. Barker, T. Cox, J. Culling, G. Naylor, and E. Porter, “Clarity: Machine Learning Challenges to Revolutionise Hearing Device Processing,” in *Forum Acusticum*, Lyon, France, Dec. 2020, pp. 3495–3497.
- [12] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, “Attention is All you Need,” in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., 2017, vol. 30.
- [13] W.-N. Hsu, A. Sriram, A. Baevski, T. Likhomanenko, Q. Xu, V. Pratap, J. Kahn, A. Lee, R. Collobert, G. Synnaeve, and M. Auli, “Robust wav2vec 2.0: Analyzing Domain Shift in Self-Supervised Pre-Training,” in *Proc. Interspeech*, 2021, pp. 721–725.
- [14] J. A. Grange and J. F. Culling, “The benefit of head orientation to speech intelligibility in noise,” *The Journal of the Acoustical Society of America*, vol. 139, no. 2, pp. 703–712, 02 2016.
- [15] F. Wilcoxon, “Individual comparisons by ranking methods,” *Biometrics Bulletin*, vol. 1, no. 6, pp. 80–83, 1945.
- [16] J. M. Kates and K. H. Arehart, “The Hearing-Aid Speech Perception Index (HASPI),” *Speech Communication*, vol. 65, pp. 75–93, 2014.
- [17] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Int. Conf. on Learning Representations (ICLR)*, San Diego, CA, USA, 2015.
- [18] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: A simple way to prevent neural networks from overfitting,” *J. of Machine Learning Res.*, vol. 15, no. 56, pp. 1929–1958, 2014.
- [19] G. Close, T. Hain, and S. Goetze, “Non intrusive intelligibility predictor for hearing impaired individuals using self supervised speech representations,” 2023.

- [20] R. Edo Zezario, F. Chen, C.-S. Fuh, H.-M. Wang, and Y. Tsao, “MBI-Net: A Non-Intrusive Multi-Branched Speech Intelligibility Prediction Model for Hearing Aids,” in *Proc. Interspeech*, 2022, pp. 3944–3948.