



Exploiting Hidden Representations from a DNN-based Speech Recogniser for Speech Intelligibility Prediction in Hearing-impaired Listeners

Zehai Tu, Ning Ma, Jon Barker

University of Sheffield, Department of Computer Science, Sheffield, UK

{ztu3, n.ma, j.p.barker}@sheffield.ac.uk

Abstract

An accurate objective speech intelligibility prediction algorithm is of great interest for many applications such as speech enhancement for hearing aids. Most algorithms measure the signal-to-noise ratios or correlations between the acoustic features of clean reference signals and degraded signals. However, these hand-picked acoustic features are usually not explicitly correlated with recognition. Meanwhile, deep neural network (DNN) based automatic speech recogniser (ASR) is approaching human performance in some speech recognition tasks. This work leverages the hidden representations from DNN-based ASR as features for speech intelligibility prediction in hearing-impaired listeners. The experiments based on a hearing aid intelligibility database show that the proposed method could make better prediction than a widely used short-time objective intelligibility (STOI) based binaural measure.

Index Terms: Intelligibility prediction, objective measures, deep neural networks, hearing-impaired listeners, hearing aids

1. Introduction

Accurate objective speech intelligibility measurement plays an important role in the development of speech enhancement, e.g. hearing aids, because subjective listening experiments can be time-consuming and expensive [1]. Since the 2010s, objective intelligibility prediction algorithms, including STOI [2], sEPSM [3], and HASPI [4], have achieved success by comparing the acoustic features of reference and degraded speech signals. However, these algorithms have been rarely evaluated on the degraded speech enhanced by data-driven DNN-based models, which have made significant progress in speech enhancement [5].

The speech recognition performance of recent DNN-based ASR models is approaching that of humans, and they have also shown similar patterns in speech recognition results [6, 7, 8, 9]. Therefore, it has been of interest to use DNN-based ASR for intelligibility prediction. Meanwhile, DNNs are naturally good feature extractors. Compared to the acoustic features proposed in the aforementioned intelligibility prediction algorithms, hidden representations of DNN-based ASR are optimised to directly correlate with recognition.

In this work, we exploit the hidden representations from one of the state-of-the-art end-to-end ASR models for intelligibility prediction on the first round Clarity Prediction Challenge (CPC1) [10], which includes a large number of binaural speech signals simulated in complex noisy environments and then enhanced by complex hearing-aid models. The results show that the similarities between ASR hidden representations of reference and processed signals could be better than the conventional acoustic features at intelligibility prediction. The results also show that using ASR hidden representations could be a better method for intelligibility prediction than simply using ASR

recognition results.

This paper is organised as follows. Section 2 reviews related objective intelligibility prediction measures. Section 3 presents the ASR model used in this work and the method for similarity computation from hidden representations. Section 4 shows the experimental setup including database and evaluation methods. The results and detailed analyses are presented and discussed in Section 5. Section 6 concludes the work and proposes future plans.

2. Related Work

Early work on intelligibility prediction for speech in additive noise is based on measured signal-to-noise ratio (SNR) in frequency bands or at modulation frequencies, e.g., the articulation index [11], the speech intelligibility index [12] and the speech transmission index [13]. However, these methods are not suitable in scenes involving (strong) non-stationary noises and non-linear processing [14]. A more recent SNR-based algorithm is sEPSM, which calculates the ratio between envelope powers of processed signals and residual noise. Similarly, GEDI [15] predicts intelligibility based on the signal-to-distortion ratio in the envelopes output from a gammachirp auditory filterbank. More recent methods have been correlation-based, with STOI being the most widely used example [16]. STOI and its variants are based on the cross-correlation between the temporal envelopes output from a one third octave filterbank of the reference and processed signals. However, they produce poor estimates for noisy speech enhanced by Wiener filtering [15] or DNN-based enhancement [17]. Additionally, HASPI incorporates a hearing loss model, and combines the measures of auditory coherence and cepstral correlation between the reference signal and the processed signal degraded by hearing loss. Most of these methods use perceptually motivated acoustic features, while not directly tuned to be correlated with speech recognition.

Previous works using ASR models for intelligibility prediction usually leverage outputs of ASR models, and most of them are evaluated on small vocabulary tasks. For example, the predicted words are used to compute the recognition correctness in [7, 8]. Alternatively, phoneme posterior probabilities from the ASR outputs are used to estimate the ASR word error rates [18] or directly predict speech recognition thresholds [19] for matrix sentence tests. Several word posterior related measures are explored in [20]. For unlimited vocabulary size, phone accuracy and phone posterior are used in [21] and [9]. To the best of our knowledge, ours is the first work to use ASR hidden representations for intelligibility prediction. The proposed method could take better advantage of the reference signals, and work particularly well in some situations, e.g., when the processed and reference signals are very different, while the ASR still makes a correct *guess*.

3. Method

In this section we describe how we leverage hidden representations from an ASR model for intelligibility prediction. ASR based on transformer architectures [22] has achieved great success recently and is used in this work. The predicted intelligibility is defined as the similarity between hidden representations from reference signals x and processed signals \hat{x} . In addition, performances of hidden representations at different levels within the ASR are investigated. Although binaural signals produced by hearing aids are studied, single-channel ASR is used for hidden representation extraction, that is, two hidden representations for left and right channels are extracted given a binaural signal. All four hidden representations from the left, right channels of reference and processed signals $H^l, H^r, \hat{H}^l, \hat{H}^r$ are used to compute the similarity.

3.1. ASR model

Figure 1 shows the architecture of the transformer-based ASR model used in this work. It consists of a convolutional neural network (CNN) based PreNet, a transformer-based encoder, and a transformer-based decoder. The PreNet is a stack of convolutional layers for better understanding global context [23]. Both the encoder and decoder are composed of a number of transformer blocks. Each encoder transformer block consists of a multi-head self-attention sub-layer and a position-wise fully connected feed-forward sub-layer. A residual connection and layer normalisation [24] are applied to both sub-layers. In the multi-head attention sub-layer, the input features are firstly mapped to query Q , key K with embedding length d_k , and value V with embedding length d_v , and the attention mechanism is computed as:

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V. \quad (1)$$

The projection and attention mechanism are run in parallel multiple times, and the concatenation of attention outputs is then multiplied by a linear projection matrix. Compared to the encoder transformer block, an extra multi-head attention sub-layer is inserted to perform the attention mechanism on the encoder output features. In addition, a positional mask is used in the decoder multi-head self-attention sub-layer to enforce that only the known previous decoded outputs are dependent.

The ASR model is optimised with the joint CTC-attention mechanism [25], i.e., a combination of Connectionist Temporal Classification (CTC) [26] and attention-based sequence-to-sequence (seq2seq) [27]. The CTC leverages repeatable intermediate label representation and a special blank label for ASR decoding, and the loss function can be expressed as:

$$\mathcal{L}_{CTC} = -\log\left(\sum_{\pi \in \beta^{-1}(l)} \prod_{m=1}^M P(z_{\pi_m}^m)\right), \quad (2)$$

where β is a function that removes repeated intermediate and blank labels, π_m is the intermediate and blank label sequence, $P(z_{\pi_m}^m)$ is the probability of π_m at time m , and l is the target label sequence. The seq2seq loss function is the sum of divergences between the ground truth label z_u and predicted token \hat{z}_u at u -th position in the transcript sequence:

$$\mathcal{L}_{seq2seq} = \sum_u P(z_u)(\log P(z_u) - \log P(\hat{z}_u)). \quad (3)$$

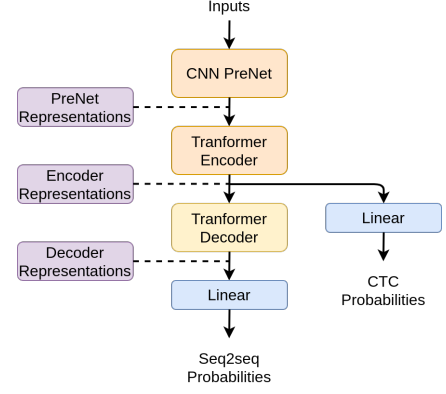


Figure 1: ASR architecture and hidden representations at three levels.

The overall loss function for ASR optimisation is:

$$\mathcal{L} = \lambda \mathcal{L}_{CTC} + (1 - \lambda) \mathcal{L}_{seq2seq}, \quad (4)$$

where λ is a predefined weighting coefficient.

3.2. Hidden representations

Three hidden representations shown in Figure 1 are studied in this work, including outputs of the CNN PreNet $H^{pre} \in \mathcal{R}^{T^{pre} \times d^{pre}}$, outputs of the transformer encoder $H^{enc} \in \mathcal{R}^{T^{enc} \times d^{enc}}$, and outputs of the transformer decoder $H^{dec} \in \mathcal{R}^{T^{dec} \times d^{dec}}$. The PreNet representations H^{pre} are viewed as low-level acoustic features. Meanwhile, the encoder representations H^{enc} can be viewed as high-level acoustic representations, as ASR models using CTC decoding does not learn a language model, and CTC output intermediate labels are independent from each other. In contrast, the seq2seq decoder is usually considered as an internal language model [28]. Therefore, the decoder representations H^{dec} are viewed as hidden representations with learnt language knowledge.

3.3. Similarity computation

The cosine similarity is used in this work as it is naturally well scaled. Given two hidden representations at a single time step, $h, \hat{h} \in \mathcal{R}^d$, the cosine similarity is computed as: $\cos(h, \hat{h}) = \frac{h \cdot \hat{h}}{\|h\| \|\hat{h}\|}$, where $\|\cdot\|$ is the $L2$ norm. For PreNet and encoder representations, the reference and processed representations of each time step are matched. The similarity at each time step ρ_t for the binaural reference and processed representations is computed from the pair of representations at this time step, i.e., $\rho_t = \cos(h_t, \hat{h}_t)$. The overall similarity between the binaural reference and processed representations is computed from the four pairs of hidden representations at each time step, i.e., $\{h^l, \hat{h}^l\}, \{h^r, \hat{h}^r\}, \{h^l, \hat{h}^r\}, \{h^r, \hat{h}^l\}$:

$$\text{sim}(H^{bi}, \hat{H}^{bi}) = \frac{1}{T} \sum_{t=1}^T \max\{\rho_t^{ll}, \rho_t^{lr}, \rho_t^{rl}, \rho_t^{rr}\}. \quad (5)$$

For decoder representations, the left and right representations of the reference and processed signals could have variable time steps, i.e., $T^l, T^r, \hat{T}^l, \hat{T}^r$ could be different. For each pair of sequences of decoder representations $\{H, \hat{H}\}$, the fast dynamic time warping algorithm [29] is applied to find the warp

path. The overall similarity of the warped pair is computed as:

$$\text{sim}(H_w, \hat{H}_w) = \frac{1}{T_w} \sum_{t=1}^{T_w} \cos(H_w(t), \hat{H}_w(t)), \quad (6)$$

where H_w and \hat{H}_w are the warped representations, T_w is the total time steps after warping. The overall binaural similarity is then computed as:

$$\text{sim}(H^{bi}, \hat{H}^{bi}) = \max \left\{ \text{sim}(H_w^l, \hat{H}_w^l), \text{sim}(H_w^r, \hat{H}_w^r), \text{sim}(H_w^r, \hat{H}_w^l), \text{sim}(H_w^l, \hat{H}_w^r) \right\}. \quad (7)$$

4. Experimental Setup

4.1. Database

CPC1 provides a large number of processed binaural speech signals by machine learning hearing-aid systems and the corresponding responses from hearing impaired listeners. Each signal represents a simulated mixture of a target speech and an interfering noise within a simulated cuboid-shaped living room, enhanced by a hearing-aid system given the audiogram (i.e., pure-tone measure of hearing thresholds at different frequencies) of a listener. Both the binaural processed signals and the corresponding anechoic reference signals are provided. The ground truth intelligibility is presented as the listener word correctness scores (WCS). A total of 6 speakers, 10 hearing aid systems and 27 listeners are included. The CPC1 includes two tracks: (1) *closed-set*, that is the listeners and systems in the evaluation set are overlapped with those in the training data; (2) *open-set*, that is the systems or listeners in the evaluation set are not included in the training data. For full details see [10]. For both tracks, the scenes in the training data are split into 70% and 30% as a training set and a development set, and the results on the evaluation set are reported.

4.2. ASR model

The SpeechBrain [30] LibriSpeech transformer ASR recipe is used in this work. 80-channel log mel-filter bank coefficients are used as input with a 25 ms window with a stride of 10 ms. The PreNet consists of three 2D convolutional layers, and the encoder and the decoder consists of twelve and six transformer blocks, respectively. The weighting coefficient λ is set 0.3 for training, and 0.4 for decoding. The dimensions at one time step for PreNet, encoder, and decoder hidden representations are 10240, 768, and 768, respectively.

4.3. Evaluation

The baseline intelligibility predictor includes the Cambridge MSBG hearing loss model [31, 32, 33, 34] and MBSTOI [16]. The MSBG hearing loss model simulates hearing abilities given a listener’s audiogram, and the MBSTOI is a refined binaural version of STOI. In addition, the correlations between WCS of listeners and of ASR models are also reported.

Three performance evaluation measures, including root mean square error (RMSE), normalised cross-correlation coefficient (NCC), and Kendall’s Tau coefficient (KT), are exploited as the evaluation metrics. As the first two metrics could be invalid for non-linear correlations, a logistic function $f(x) = 1/[1 + \exp(ax + b)]$ is applied to the predicted intelligibility

Table 1: Evaluation results on the CPC1 of various methods.

	RMSE ↓	NCC ↑	KT ↑
Closed-set			
Baseline	0.285	0.621	0.398
ASR WCS	0.250	0.729	0.523
PreNet representations	0.347	0.299	0.182
Encoder representations	0.237	0.758	0.487
Decoder representations	0.231	0.773	0.498
Open-set			
Baseline	0.365	0.529	0.391
ASR WCS	0.250	0.723	0.534
PreNet representations	0.356	0.254	0.136
Encoder representations	0.241	0.751	0.534
Decoder representations	0.235	0.763	0.530

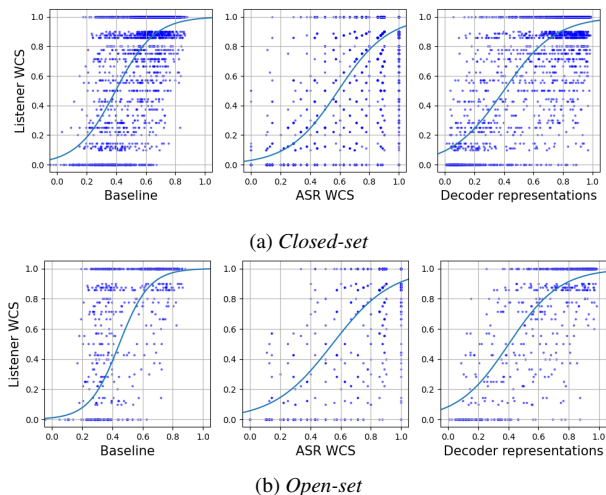


Figure 2: Listener WCS versus predicted intelligibility by various methods and the corresponding logistic mapping functions.

to examine the monotonic relation, following the conventions of previous works, including [2, 16]. For the proposed method, the parameters a and b are optimised on the development set. For the baseline system, the parameters are optimised on all the CPC1 training data, as described in [10].

5. Experiments and Results

5.1. Hidden representations

The training of the default ASR model starts from the pre-trained model on the LibriSpeech¹ (LS). Therefore, it has a strong knowledge on clean speech. Furthermore, it is optimised with LS *train-clean-100* set added with noises from the training set in the first round Clarity Enhancement Challenge [35] (CLS) for ten epochs. Finally, the ASR model is optimised on CPC1 training set for another ten epochs. In addition, the MSBG hearing loss model is used to process the signals when training and testing on CPC1. The correlations between the listener WCS and the baseline (MSBG+MBSTOI) prediction, ASR WCS, predicted intelligibility with different ASR hidden representations, are shown in Table 1. Figure 2 also shows the listener WCS against the predicted intelligibility by the base-

¹huggingface.co/speechbrain/asr-transformer-transformerlm-librispeech

Table 2: Evaluation results on the closed-set of decoder representations from different ASR models.

MSBG	Training data	RMSE ↓	NCC ↑	KT ↑
with	LS	0.264	0.692	0.449
	LS+CLS	0.243	0.746	0.464
	LS+CPC1	0.233	0.768	0.503
	LS+CLS+CPC1	0.231	0.773	0.498
w/o	LS+CLS+CPC1	0.234	0.767	0.476

Table 3: Listener- and system-wise evaluation results on the closed-set of predicted intelligibility.

	RMSE ↓	NCC ↑	KT ↑
Listener-wise			
Baseline	0.078	0.414	0.311
Decoder representations	0.078	0.419	0.407
System-wise			
Baseline	0.147	0.798	0.244
Decoder representations	0.048	0.982	0.644

line, ASR WCS, the decoder representations, and their corresponding logistic mapping functions. Both the results of the *closed-set* and *open-set* indicate that the similarity between the reference and processed high-level hidden representations could outperform the baseline and ASR WCS at intelligibility prediction in terms of RMSE and NCC. The ASR WCS predictions are advantageous with regard to KT because WCS is discrete, i.e., in which case *tied* pairs are more likely to appear. Between the two high-level hidden representations, the decoder ones including language model knowledge are better than the encoder ones which represent high-level acoustic features in terms of RMSE and NCC, while the KT scores are close. The following experiments and analyses are conducted on the *closed-set*.

5.2. Data mismatch

For the purpose of investigating the influence of data mismatch (i.e., different distribution of training and evaluation data) on ASR models for intelligibility prediction, four different ASR models with different training data knowledge (LS, LS+CLS, LS+CPC1, LS+CLS+CPC1) are probed. The MSBG model is used for all models as preprocessing for hearing loss simulation. ASRs trained on CLS can be considered to have knowledge of noisy speech, and those trained on CPC1 can be considered to have knowledge of processed noisy speech by hearing-aid systems. The correlations between the predicted intelligibility with decoder representations and the ground truth WCS are shown in Table 2. The results show that the ASR models trained with CPC1 training data (LS+CPC1, LS+CLS+CPC1) could make optimal predictions, while the latter one is slightly better in terms of RMSE and NCC because it has knowledge of noisy speech. Meanwhile, the ASR models with no knowledge of CPC1 data (LS, LS+CLS) could also achieve competitive results. It is worth noting that the ASR model trained only on clean LS signals could still outperform the baseline system.


5.3. MSBG hearing loss model

The influence of the MSBG hearing loss model is also investigated. The intelligibility prediction results of ASR models trained on LS+CLS+CPC1 with and without the MSBG model for hearing loss simulation are also shown in Table 2. The results indicate that the MSBG hearing loss model could offer a


slight advantage on intelligibility prediction for the ASR hidden representations.

5.4. Listener- and system-wise correlation

The results of the listening experiments provided by CPC1 can be noisy, because of the not strictly constrained speech materials, the large size vocabulary, etc. Therefore, both the listener WCS and the predicted intelligibility scores are averaged on listeners and hearing-aid systems for more conclusive analysis. The average listener WCS, the predicted intelligibility from the baseline, and the proposed decoder representation similarity from the ASR trained on LS+CLS+CPC1 with the MSBG hearing loss model on different listeners and hearing-aid systems with their corresponding error bars are shown in Figure 3a and Figure 3b. The listener- and system-wise evaluation results on the *closed-set* are measured and shown in Table 3. The intelligibility predicted from the ASR decoder representations could gain a slight advantage over the baseline on listener-wise intelligibility scores. For system-wise intelligibility, the proposed method is much more correlated with the listener WCS than the baseline.



(a) Listener-wise



(b) System-wise

Figure 3: Listener- and system-wise average intelligibility with standard errors on the closed-set.

6. Conclusions

This paper has proposed a novel method for speech intelligibility prediction by leveraging DNN-based ASR hidden representations. The similarity of the hidden representations from an ASR neural network is measured between a clean reference signal and the corresponding processed signal which is used for intelligibility prediction. The experimental results on the recent CPC1 database, which provides machine learning hearing aid processed binaural signals, have shown that the proposed method can outperform other intrusive methods such as MB-STOI, the refined binaural variant of STOI. Future work will focus on expanding the evaluation of the proposed method to different databases, such as [6], with more baseline methods, including STOI and HASPI.

7. References

- [1] T. H. Falk, V. Parsa, J. F. Santos, K. Arehart, O. Hazrati, R. Huber, J. M. Kates, and S. Scollie, "Objective quality and intelligibility prediction for users of assistive listening devices: Advantages and limitations of existing tools," *IEEE signal processing magazine*, vol. 32, no. 2, pp. 114–124, 2015.
- [2] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, "An algorithm for intelligibility prediction of time–frequency weighted noisy speech," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 7, pp. 2125–2136, 2011.
- [3] S. Jørgensen and T. Dau, "Predicting speech intelligibility based on the signal-to-noise envelope power ratio after modulation-frequency selective processing," *The Journal of the Acoustical Society of America*, vol. 130, no. 3, pp. 1475–1487, 2011.
- [4] J. M. Kates and K. H. Arehart, "The hearing-aid speech perception index (haspi)," *Speech Communication*, vol. 65, pp. 75–93, 2014.
- [5] D. Wang and J. Chen, "Supervised speech separation based on deep learning: An overview," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 10, pp. 1702–1726, 2018.
- [6] J. Barker and M. Cooke, "Modelling speaker intelligibility in noise," *Speech Communication*, vol. 49, no. 5, pp. 402–417, 2007.
- [7] M. R. Schädler, A. Warzybok, S. Hochmuth, and B. Kollmeier, "Matrix sentence intelligibility prediction using an automatic speech recognition system," *International Journal of Audiology*, vol. 54, no. sup2, pp. 100–107, 2015.
- [8] L. Fontan, I. Ferrané, J. Farinas, J. Pinquier, J. Tardieu, C. Magne, P. Gaillard, X. Aumont, and C. Füllgrabe, "Automatic speech recognition predicts speech intelligibility and comprehension for listeners with simulated age-related hearing loss," *Journal of Speech, Language, and Hearing Research*, vol. 60, no. 9, pp. 2394–2405, 2017.
- [9] K. Arai, S. Araki, A. Ogawa, K. Kinoshita, T. Nakatani, and T. Irino, "Predicting intelligibility of enhanced speech using posteriors derived from dnn-based asr system," in *INTERSPEECH*, 2020, pp. 1156–1160.
- [10] J. Barker, M. Akeroyd, T. J. Cox, J. F. Culling, J. Firth, S. Graetzer, H. Griffiths, L. Harris, G. Naylor, Z. Podwinska, E. Porter, R. Viveros Munoz *et al.*, "The 1st clarity prediction challenge: A machine learning challenge for hearing aid intelligibility prediction," in *INTERSPEECH*, 2022.
- [11] K. D. Kryter, "Methods for the calculation and use of the articulation index," *The Journal of the Acoustical Society of America*, vol. 34, no. 11, pp. 1689–1697, 1962.
- [12] "Methods for calculation of the speech intelligibility index," *S3.5-1997, ANSI*, 1997.
- [13] H. J. Steeneken and T. Houtgast, "A physical method for measuring speech-transmission quality," *The Journal of the Acoustical Society of America*, vol. 67, no. 1, pp. 318–326, 1980.
- [14] S. Graetzer and C. Hopkins, "Intelligibility prediction for speech mixed with white gaussian noise at low signal-to-noise ratios," *The Journal of the Acoustical Society of America*, vol. 149, no. 2, pp. 1346–1362, 2021.
- [15] K. Yamamoto, T. Irino, T. Matsui, S. Araki, K. Kinoshita, and T. Nakatani, "Predicting speech intelligibility using a gammachirp envelope distortion index based on the signal-to-distortion ratio," in *INTERSPEECH*, 2017, pp. 2949–2953.
- [16] A. H. Andersen, J. M. de Haan, Z.-H. Tan, and J. Jensen, "Refinement and validation of the binaural short time objective intelligibility measure for spatially diverse conditions," *Speech Communication*, vol. 102, pp. 1–13, 2018.
- [17] F. B. Gelderblom, T. V. Tronstad, and E. M. Viggen, "Subjective intelligibility of deep neural network-based speech enhancement," in *INTERSPEECH*, 2017, pp. 1968–1972.
- [18] A. M. C. Martinez, C. Spille, J. Roßbach, B. Kollmeier, and B. T. Meyer, "Prediction of speech intelligibility with dnn-based performance measures," *Computer Speech & Language*, vol. 74, p. 101329, 2021.
- [19] J. Roßbach, S. Röttges, C. F. Hauth, T. Brand, and B. T. Meyer, "Non-intrusive binaural prediction of speech intelligibility based on phoneme classification," in *ICASSP*, 2021, pp. 396–400.
- [20] M. Karbasi, S. Bleeck, and D. Kolossa, "Non-intrusive speech intelligibility prediction using automatic speech recognition derived measures," *arXiv preprint arXiv:2010.08574*, 2020.
- [21] K. Arai, S. Araki, A. Ogawa, K. Kinoshita, T. Nakatani, K. Yamamoto, and T. Irino, "Predicting speech intelligibility of enhanced speech using phone accuracy of dnn-based asr system," in *INTERSPEECH*, 2019, pp. 4275–4279.
- [22] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *NeurIPS*, 2017, pp. 5998–6008.
- [23] W. Han, Z. Zhang, Y. Zhang, J. Yu, C.-C. Chiu, J. Qin, A. Gulati, R. Pang, and Y. Wu, "ContextNet: Improving convolutional neural networks for automatic speech recognition with global context," in *INTERSPEECH*, 2020, pp. 3610–3614.
- [24] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.
- [25] S. Kim, T. Hori, and S. Watanabe, "Joint CTC-attention based end-to-end speech recognition using multi-task learning," in *ICASSP*. IEEE, 2017, pp. 4835–4839.
- [26] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 369–376.
- [27] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, "Attention-based models for speech recognition," in *NeurIPS*, 2017.
- [28] Z. Meng, S. Parthasarathy, E. Sun, Y. Gaur, N. Kanda, L. Lu, X. Chen, R. Zhao, J. Li, and Y. Gong, "Internal language model estimation for domain-adaptive end-to-end speech recognition," in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 243–250.
- [29] S. Salvador and P. Chan, "Toward accurate dynamic time warping in linear time and space," *Intelligent Data Analysis*, vol. 11, no. 5, pp. 561–580, 2007.
- [30] M. Ravanelli, T. Parcollet, P. Plantinga, A. Rouhe, S. Cornell, L. Lugosch, C. Subakan, N. Dawalatabad, A. Heba, J. Zhong *et al.*, "SpeechBrain: A general-purpose speech toolkit," *arXiv preprint arXiv:2106.04624*, 2021.
- [31] T. Baer and B. C. J. Moore, "Effects of spectral smearing on the intelligibility of sentences in noise," *JASA*, vol. 94, no. 3, pp. 1229–1241, 1993.
- [32] —, "Effects of spectral smearing on the intelligibility of sentences in the presence of interfering speech," *JASA*, vol. 95, no. 4, pp. 2277–2280, 1994.
- [33] B. C. J. Moore and B. R. Glasberg, "Simulation of the effects of loudness recruitment and threshold elevation on the intelligibility of speech in quiet and in a background of speech," *JASA*, vol. 94, no. 4, pp. 2050–2062, 1993.
- [34] M. A. Stone and B. C. Moore, "Tolerable hearing aid delays. i. estimation of limits imposed by the auditory path alone using simulated hearing losses," *Ear and Hearing*, vol. 20, no. 3, pp. 182–192, 1999.
- [35] S. Graetzer, J. Barker, T. J. Cox, M. Akeroyd, J. F. Culling, G. Naylor, E. Porter, R. Viveros Munoz *et al.*, "Clarity-2021 challenges: Machine learning challenges for advancing hearing aid processing," in *INTERSPEECH*, 2021, pp. 686–690.