



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/208186/>

Version: Accepted Version

---

**Proceedings Paper:**

Mogridge, R., Close, G., Sutherland, R. et al. (2024) Non-intrusive speech intelligibility prediction for hearing-impaired users using intermediate ASR features and human memory models. In: ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2024), 14-19 Apr 2024, Seoul, Korea. Institute of Electrical and Electronics Engineers (IEEE), pp. 306-310. ISBN: 979-8-3503-4486-8. ISSN: 1520-6149. EISSN: 2379-190X.

<https://doi.org/10.1109/ICASSP48485.2024.10447597>

---

© 2024 The Author(s). Except as otherwise noted, this author-accepted version of a conference paper published in publication in International Conference on Acoustics, Speech, and Signal Processing (ICASSP) is made available via the University of Sheffield Research Publications and Copyright Policy under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

**Reuse**

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.

# NON-INTRUSIVE SPEECH INTELLIGIBILITY PREDICTION FOR HEARING-IMPAIRED USERS USING INTERMEDIATE ASR FEATURES AND HUMAN MEMORY MODELS

*Rhiannon Mogridge, George Close, Robert Sutherland, Thomas Hain, Jon Barker,  
Stefan Goetze, Anton Ragni*

Speech and Hearing Group, Dept. of Computer Science, University of Sheffield, UK

## ABSTRACT

Neural networks have been successfully used for non-intrusive speech intelligibility prediction. Recently, the use of feature representations sourced from intermediate layers of pre-trained self-supervised and weakly-supervised models has been found to be particularly useful for this task. This work combines the use of Whisper ASR decoder layer representations as neural network input features with an exemplar-based, psychologically motivated model of human memory to predict human intelligibility ratings for hearing-aid users. Substantial performance improvement over an established intrusive HASPI baseline system is found, including on enhancement systems and listeners unseen in the training data, with a root mean squared error of 25.3 compared with the baseline of 28.7.

*Index Terms*— speech recognition, intelligibility prediction, hearing impairment

## 1. INTRODUCTION

Hearing loss is a widespread problem that affects approximately 466 million people worldwide (around 6% of the world population), though this problem is only expected to grow; by 2030 it is predicted to impact 630 million people worldwide [1]. Age correlates with the chance that a person will be affected by hearing impairment [2], and the population is expected to age. From 2015 to 2050, the proportion of the population aged over 60 is expected to almost double, rising from 12% to 22% [3]. As hearing impairment typically worsens in an individual, their ability to make intelligible the speech that they hear decreases.

Successful development of hearing aid (HA) technology [4, 5] requires assessment, which is expensive and time-consuming when conducted by human listeners [6, 7]. Automated intelligibility metrics, which are designed to mimic human assessment, are more cost-effective and can also be used as training objectives [8, 9]. Developing and improving estimators for such metrics is therefore essential.

The Clarity Prediction Challenge 2 (CPC2) [10] builds on the prior Clarity Prediction Challenge 1 (CPC1) [11]. It provides a comparatively large dataset of noisy audio processed by hearing aid systems, each with an associated intelligibility score obtained from listening tests with hearing-impaired human listeners.

The challenge task is to predict the intelligibility score of hearing-impaired listeners given the audio and some additional information such as a representation of the hearing loss of the listener. The challenge has both a non-intrusive track, where only

the *noisy* signal processed by the hearing aid can be used, and an intrusive track, where a *clean* version of the input audio can also be used.

In this work, an intelligibility estimator for the non-intrusive challenge task is proposed, which makes use of a feature transformation of the audio signal, derived from intermediate representations of a pre-trained automatic speech recognition (ASR) model (cf. Section 3.1). This is combined with a model of human memory, which has its origins in the field of human psychology, to predict the intelligibility score of hearing aid output audio (cf. Section 3.2.2).

The remainder of this work is organised as follows: Section 2 briefly describes the CPC2, including the dataset and baseline model. Section 3 covers the proposed features and model architecture. Results are presented in Section 4 and Section 5 concludes the paper.

## 2. CLARITY PREDICTION CHALLENGE 2

The Clarity Project runs two challenges in sequence with the aim of improving HA technology: the Clarity Enhancement Challenge (CEC) and the Clarity Prediction Challenge (CPC). The CEC objective is to design systems to enhance noisy signals for hearing-impaired listeners. The CPC objective is to assess the intelligibility of the CEC systems.

### 2.1. Challenge data

The CPC2 data consists of tuples of a speech signal  $\hat{s}[n]$  and its corresponding correctness value  $i$ , obtained from listening tests with hearing-impaired listeners.

The signal  $\hat{s}[n]$  is the enhanced outputs of hearing aid systems with binaural input  $x[n]$ , being an artificially corrupted version of clean reference audio  $s[n]$  with additive noise  $v[n]$ . The correctness value  $i$  is the percentage of words which a hearing-impaired listener was able to correctly reproduce from the speech signal  $\hat{s}[n]$  they listened to. The challenge data also contains additional information such as left/right ear’s representations of the listeners’ hearing loss as audiograms  $\mathbf{a}_l$  and  $\mathbf{a}_r$ . All audio signals are stereo with a left and right channel.

The resulting data is partitioned into three train sets, each paired with a disjoint evaluation set. Each evaluation set covers listeners and hearing aid enhancement systems which are unseen in its corresponding training set, meaning that prediction models need to generalise to unseen listeners and systems. Set 1 has a training set of size 8599 and an evaluation set of size 305 audio samples. Set 2 has a training set of size 8135 and an evaluation set of size 294. Set 3 has a training set of size 7896 and an evaluation set of size 298. There are around 40 hours of audio in total. Audio with very low intelligibility (correctness 0) and very high intelligibility (correctness 100) are over-represented, as shown in Figure 1.

---

This work was supported by the Centre for Doctoral Training in Speech and Language Technologies (SLT) and their Applications funded by UK Research and Innovation [grant number EP/S023062/1]. This work was also supported by Toshiba and WS Audioogy.




Fig. 1. Distribution of true correctness values in the training data.

## 2.2. Prior approaches

A number of different approaches were taken in the first iteration of the challenge in CPC1 [11]. The best performing non-intrusive approach [12] uses an uncertainty measure derived from state-of-the-art ASR systems as a proxy for human intelligibility, finding a strong correlation between the two measures. Other successful approaches [13, 14] make use of powerful feature representations derived from self-supervised speech representations (SSSRs) as inputs to neural speech intelligibility prediction models, while others use neural network structures which have been shown to be useful in the related task of human speech quality rating prediction [15].

CPC2 differs from CPC1 in that its evaluation sets are disjoint in terms of listener and hearing aid system relative to its training sets. This means that some of the better-performing approaches to CPC1, which operated effectively as predictors of the hearing aid system, were not at all useful in CPC2, which was discovered in early experiments for this work with the CPC2 data. As such, our proposed system for CPC2 builds on the best-performing approaches to CPC1, while ensuring that it can generalise to unseen data.

## 2.3. Challenge baseline

The baseline system provided by the challenge organisers [10] makes use of the Hearing-Aid Speech Perception Index (HASPI), version 2 [16]. This is an intrusive system that makes use of both the enhanced noisy signal  $\hat{s}[n]$  and the clean speech signal  $s[n]$ . A HASPI score is computed for both the left and right ear signals and logistic regression is used to predict the correctness scores from the higher HASPI score.

## 3. SYSTEM ARCHITECTURE

This section describes the proposed approach to the CPC2 task. The approach consists of neural networks which take a recent ASR derived representation of the hearing aid output signal  $\hat{s}[n]$  as input and return a prediction  $\hat{i}$  of the correctness value.

### 3.1. Features

The Whisper model [17] is an ASR model pre-trained on 680,000 hours of multi-lingual data for tasks including English and non-English speech transcription and voice activity detection. The model architecture is an encoder-decoder transformer [18], with 12 encoder layers and 12 decoder layers for the *small* Whisper model<sup>1</sup> used in this work.

<sup>1</sup><https://huggingface.co/openai/whisper-small>

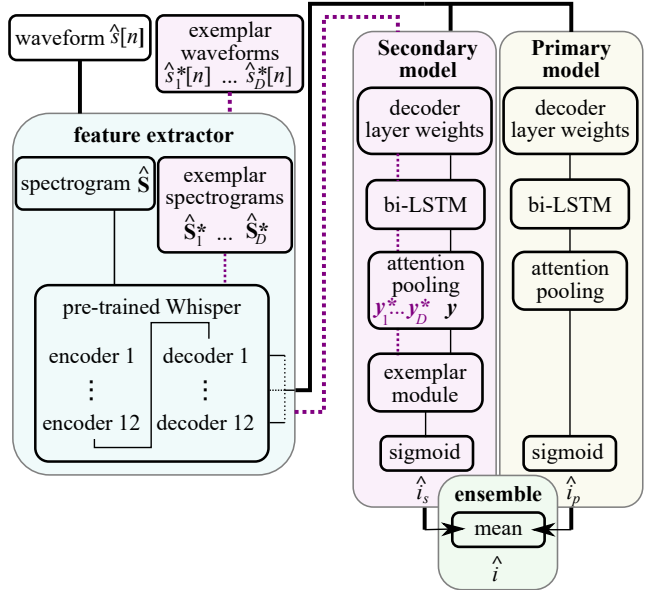


Fig. 2. Model architecture of proposed SI prediction.

Given the enhanced speech signal  $\hat{s}[n]$ , with corresponding spectrogram representation  $\hat{S}$ , the input to the proposed intelligibility prediction model is the output of the 12 decoder layers of the Whisper model. The Whisper decoder layers output represents word-level features, with dimension  $W \times 768 \times 12$ , where  $W$  is the predicted number of words in the utterance, 768 is the feature dimension of each decoder layer for the *small* Whisper model, and 12 is the number of decoder layers. Note that while  $W$  varies from utterance to utterance, for a given utterance it will remain fixed through the decoder layers. The parameters of the Whisper model are frozen during training of the metric prediction model described below, i.e., the Whisper model is used as a feature transform.

### 3.2. Model structure

An ensemble of two models for speech intelligibility (SI) prediction is used, as shown in Figure 2, and the results of these two models are combined by averaging.

A model structure following work on the CPC1 in [14] is chosen for the *primary* SI prediction network (cf. Section 3.2.1), depicted to the right in Figure 2. This model structure has previously been successfully applied to the task of human quality label prediction [19] using pre-trained SSSR representations as its input feature. This approach was also found to be useful for the CPC1 task [14], however generalization to unseen hearing aid systems was poor.

The *secondary* SI prediction network incorporates an exemplar-informed module based on a simplified theory of human memory [20] (cf. Section 3.2.2), and is shown in the middle-right of Figure 2. Humans are believed to make use of specific examples, or exemplars, for memory-based tasks [21, 22, 23, 24]. Humans are also able to non-intrusively assess speech signals, i.e., without direct reference. Since the challenge objective is to predict human responses, incorporating an exemplar-informed component may provide benefits or insight.

The output of the ensemble  $\hat{i}$  for a given input signal  $\hat{s}[n]$  is the mean of the outputs of the primary  $\hat{i}_p$  and secondary systems  $\hat{i}_s$ .

### 3.2.1. Primary SI prediction model

The model structure uses a learnable weighted sum of the Whisper representations, implemented as a learnable linear layer with 12 parameters, all initialised to 1, followed by a softmax to ensure that the layer weights sum to 1. This representation of dimension  $W \times 768$  is then processed by 2 bidirectional long short-term memory (BLSTM) layers with an input size of 768 and a hidden layer size of 384. Finally, an attention pooling feed-forward layer with sigmoid activation outputs to a single neuron which represents the primary predicted correctness value  $i_p$  normalized between 0 and 1. The primary model has approximately 8.3 M parameters.

### 3.2.2. Secondary SI prediction model: exemplar-informed

The secondary model differs from the primary model in that the attention-pooling output feeds into an exemplar-informed module based on a simplified theory of human memory [20]. The exemplar-informed module incorporates a memory set of  $D$  exemplars, which are speech signals  $\hat{s}_1^*[n], \dots, \hat{s}_D^*[n]$  drawn from the training data, with corresponding correctness values  $i_1^*, \dots, i_D^*$ . The exemplars can be changed during training and for inference. Let  $\mathbf{y}$  be the output of the attention pooling for input  $\hat{s}[n]$  (see Figure 2). The exemplars,  $\hat{s}_1^*[n], \dots, \hat{s}_D^*[n]$ , are processed in the same way as the input, producing exemplar outputs from the attention pooling  $\mathbf{y}_1^*, \dots, \mathbf{y}_D^*$ . The output,  $r$ , of the exemplar module is given by

$$a = \sum_{d=1}^D \frac{\mathbf{f}(\mathbf{y}) \cdot \mathbf{g}(\mathbf{y}_d^*)}{\|\mathbf{f}(\mathbf{y})\| \|\mathbf{g}(\mathbf{y}_d^*)\|} i_d^* \quad (1)$$

$$r = h(a) \quad (2)$$

The functions  $\mathbf{f} : \mathbb{R}^{768} \rightarrow \mathbb{R}^{768}$ ,  $\mathbf{g} : \mathbb{R}^{768} \rightarrow \mathbb{R}^{768}$  and  $h : \mathbb{R} \rightarrow \mathbb{R}$  are all learned affine transformations. The value  $a$  is a combination of the exemplar labels, weighted by their similarity to the input. This passes through a single linear neuron to produce  $r$ , and then through a sigmoid activation, which yields the secondary model’s prediction,  $\hat{i}_s$ , normalised to fall between 0 and 1. The exemplar model has approximately 10 M parameters.

### 3.3. Experimental setup

All audio used as input to Whisper is re-sampled to 16kHz and padded/truncated to be 30 seconds long. In the case of the CPC2 data, all recordings were shorter than 30 seconds, so were padded, and downsampled from 32kHz to 16kHz. From this time-domain signal, the 80 channel log magnitude Mel spectrogram is computed, using a window of 25ms and a stride of 10ms. During training, both left and right ear signals are used as independent samples with the same correctness label. During inference, the signal that produces the highest correctness value is used to account for the better-ear effect [25].

For each of the three splits, two listeners and two systems were randomly selected to form a disjoint validation set. All data with these listeners and systems were removed from the training set. A randomly selected non-disjoint validation set consisting of 10% of the remaining training data was also formed. The majority of model selection and hyperparameter tuning was performed using these validation sets, to test generalisation to unseen listeners and systems. For the final models, the disjoint validation set and all listeners/systems associated with it were merged back into the training data to make the best use of resources.

The primary and secondary models are trained separately with mean squared error loss. The primary model is trained for 25 epochs with batch size 8, learning rate  $10^{-5}$  and weight decay  $10^{-4}$ . The secondary model is trained for 50 epochs with learning rate  $2 \times 10^{-6}$  and weight decay  $10^{-4}$ . During training and validation,  $D = 8$  exemplars are chosen randomly from the training data for each mini-batch.

## 4. RESULTS AND DISCUSSION

Table 1 shows the results for the primary, secondary and ensemble models on each of the data splits, as well as the overall result for the CPC2 baseline. The validation sets contain enhancement systems and listeners seen in the corresponding training data. The evaluation sets are disjoint, containing only unseen enhancement systems and listeners. The primary, secondary and ensemble prediction networks all beat the baseline on all evaluation sets.


**Table 1.** Validation and evaluation set results.

Model	RMSE								
	validation split				evaluation split				
	1	2	3	all	1	2	3	all	
CPC2 baseline									<b>28.7</b>
Primary	21.6	23.5	22.8	<b>22.7</b>	28.2	23.8	23.3		<b>25.3</b>
Secondary	21.7	23.5	22.7	<b>22.7</b>	29.1	24.5	23.4		<b>25.8</b>
Ensemble	21.6	23.4	22.7	<b>22.5</b>	28.6	23.9	23.2		<b>25.3</b>

### 4.1. Primary and secondary models

The primary and secondary systems show similar performance on the validation sets, with the ensemble of the two outperforming either. This is not replicated on the evaluation set, where the primary model outperforms the secondary model, and performs equivalently to the ensemble.

### 4.2. Performance by intelligibility



**Fig. 3.** Model performance by true correctness.

Figure 3 shows the performance of the ensemble model for different correctness values. The model performs well for very low intelligibility (0 correctness) and for very high intelligibility (100 correctness) but performs less well between the two extremes. This corresponds with the distribution of true correctness scores in the training data (see Figure 1), in which 0 and 100 correctness are over-represented.

### 4.3. Model performance on unseen enhancement systems

All the models show lower performance on Evaluation Set 1. This appears to stem from the presence of audio enhanced by enhance-




Fig. 4. Model performance by mean hearing aid system correctness.

ment system E001 (the baseline in Clarity Enhancement Challenge 2) in this evaluation set. Audio enhanced by this system has an average correctness value of 28.7%, which is significantly lower than the average of the other two enhancement systems in Evaluation Set 1, E022 and E031, which have average correctness values of 73.0% and 84.3%, respectively.

Although the proposed model can generalise to unseen enhancement systems, it predicts correctness less accurately on enhancement systems with lower performance. Figure 4 shows the proposed model’s performance by enhancement system correctness across all enhancement systems. There is a clear trend, with the proposed system more accurately predicting the correctness for better-performing enhancement systems which produce outputs with high correctness ratings. Conversely, the proposed system less accurately predicts outputs from poorly-performing enhancement systems which produce outputs with low correctness ratings. Figure 5 shows the predicted (left) and true (right) correctness across all enhancement systems, showing the proposed model overestimates the correctness of the two worst-performing enhancement systems, E001 and E038, while generally slightly underestimating the correctness of the better-performing enhancement systems.

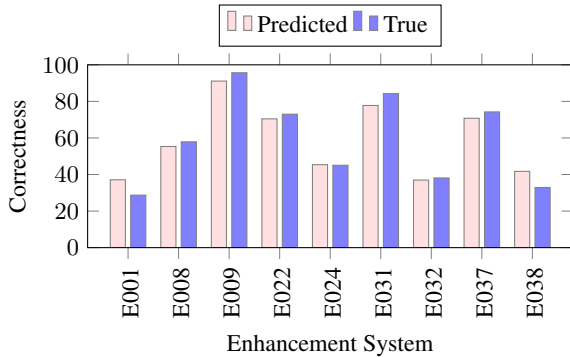


Fig. 5. Performance of proposed prediction system depending on enhancement system

#### 4.4. Whisper layer weights

The learned weights for the Whisper decoder layers from the primary model are shown in Figure 6. These show how the model used the information to weigh each decoder layer feature, and therefore the higher the value the more useful the model finds the layer to be. The weights are shown for each of the three models trained on the different training splits.

The general pattern for the different training splits is similar, with layers 7 and 8 having the highest weights across all splits. This suggests that layers 7 and 8 contain the most relevant information for intelligibility. Interestingly, the model trained on Split 3 learns weights that emphasise layers 7 and 8 more strongly.

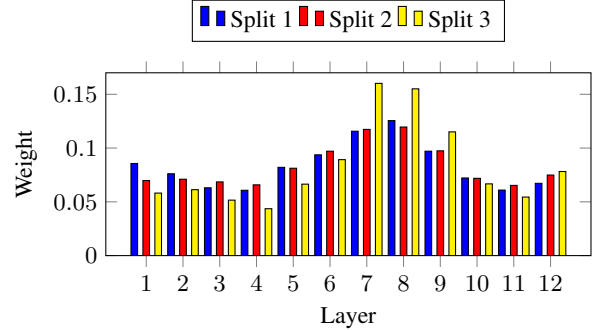


Fig. 6. Learned weights for the primary model Whisper decoder layers.

#### 4.5. Performance comparison

Figure 7 shows the performance of the proposed model (P002) compared to all other challenge entries, as well as the challenge HASPI baseline. *Prior* is a system which always predicts the average of the intelligibility over the training set, regardless of the input. The proposed system is outperformed by only one other system, P011, which also utilises Whisper-derived features. The difference in performance is very slight, with P011 achieving an RMSE score of 25.1, while our system achieves 25.3.




Fig. 7. Comparison of performance of CPC2 entries. A \* denotes an intrusive system.

## 5. CONCLUSIONS

We show that Whisper decoder layers are a useful feature representation for speech intelligibility prediction, with layers 7 and 8 appearing to be the most relevant.

Our proposed system performs substantially better than the HASPI regression baseline and all but one of the other challenge approaches, even outperforming intrusive systems which had access to the clean reference signal. It is able to generalise to unseen enhancement systems and listeners.

## 6. REFERENCES

- [1] World Health Organization, “Addressing the Rising Prevalence of Hearing Loss,” 2018, ISBN: 9789241550260.
- [2] J. RENNIES, S. GOETZE, and J.-E. APPELL, “Personalized Acoustic Interfaces for Human-Computer Interaction,” in *Human-Centered Design of E-Health Technologies: Concepts, Methods and Applications*, M. Ziefle and C. Röcker, Eds., chapter 8, pp. 180–207. IGI Global, 2011.
- [3] World Health Organisation, “Ageing and Health,” <https://www.who.int/news-room/fact-sheets/detail/ageing-and-health>, Accessed: 2023-07-26.
- [4] S. Doclo, W. Kellermann, S. Makino, and S. E. Nordholm, “Multichannel Signal Enhancement Algorithms for Assisted Listening Devices: Exploiting spatial diversity using multiple microphones,” *IEEE Signal Processing Magazine*, vol. 32, no. 2, pp. 18–30, 2015.
- [5] S. Goetze, F. Xiong, J. RENNIES, T. Rohdenburg, and J. Appell, “Hands-Free Telecommunication for Elderly Persons Suffering from Hearing Deficiencies,” in *IEEE Int. Conf. on E-Health Networking, Application and Services (Healthcom’10)*, 2010.
- [6] 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.
- [7] A. Warzybok, I. Kodrasi, J. Jungmann, E. Habets, T. Gerkmann, A. Mertins, S. Doclo, B. Kollmeier, and S. Goetze, “Subjective Speech Quality and Speech Intelligibility Evaluation of Single-Channel Dereverberation Algorithms,” in *Int. Workshop on Acoustic Signal Enhancement (IWAENC 2014)*, France, Sep. 2014.
- [8] C. Völker, A. Warzybok, and S. Ernst, “Comparing Binaural Pre-processing Strategies III: Speech Intelligibility of Normal-Hearing and Hearing-Impaired Listeners,” *Trends in Hearing*, vol. 19, 2015.
- [9] B. Cauchi, K. Siedenburg, J. F. Santos, T. H. Falk, S. Doclo, and S. Goetze, “Non-Intrusive Speech Quality Prediction Using Modulation Energies and LSTM-Network,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 7, pp. 1151–1163, July 2019.
- [10] J. Barker, M. Akeroyd, W. Bailey, T. J. Cox, J. F. Culling, J. Firth, S. Graetzer, and G. Naylor, “The 2nd Clarity Prediction Challenge: A machine learning challenge for hearing aid intelligibility prediction,” in *ICASSP*, 2024.
- [11] J. Barker, M. Akeroyd, T. J. Cox, J. F. Culling, J. Firth, S. Graetzer, H. Griffiths, L. Harris, G. Naylor, Z. Podwinska, E. Porter, and R. V. Munoz, “The 1st Clarity Prediction Challenge: A machine learning challenge for hearing aid intelligibility prediction,” in *Proc. Interspeech*, 2022, pp. 3508–3512.
- [12] Z. Tu, N. Ma, and J. Barker, “Unsupervised Uncertainty Measures of Automatic Speech Recognition for Non-intrusive Speech Intelligibility Prediction,” in *Proc. Interspeech*, 2022, pp. 3493–3497.
- [13] R. E. 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,” 2022.
- [14] G. Close, T. Hain, and S. Goetze, “Non Intrusive Intelligibility Predictor for Hearing Impaired Individuals using Self Supervised Speech Representations,” in *Proc. Workshop on Speech Foundation Models and their Performance Benchmarks (SPARKS), ASRU satellite workshop*, Taipei, Taiwan, 2023.
- [15] G. Close, S. Hollands, T. Hain, and S. Goetze, “Non-intrusive Speech Intelligibility Metric Prediction for Hearing Impaired Individuals,” in *Proc. Interspeech*, 2022, pp. 3483–3487.
- [16] J. M. Kates and K. H. Arehart, “The Hearing-aid Speech Perception Index (HASPI) Version 2,” *Speech Communication*, vol. 131, pp. 35–46, 2021.
- [17] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust Speech Recognition via Large-Scale Weak Supervision,” 2022.
- [18] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is All You Need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [19] B. Tamm, H. Balabin, R. Vandenberghe, and H. V. hamme, “Pre-trained Speech Representations as Feature Extractors for Speech Quality Assessment in Online Conferencing Applications,” in *Interspeech*. Sep 2022, ISCA.
- [20] D. Hintzman, “MINERVA 2: a Simulation Model of Human Memory,” *Behaviour Research Methods, Instruments & Computers*, vol. 16, pp. 96–101, 03 1984.
- [21] R. M. Nosofsky, T. J. Palmeri, and M. Stephen C, “Rule-Plus-Exception Model of Classification Learning,” *Psychological Review*, vol. 101, no. 1, pp. 53–79, 1994.
- [22] M. A. Erickson and J. K. Kruschke, “Rules and Exemplars in Category Learning,” *Journal of Experimental Psychology: General*, vol. 127, 1998.
- [23] J. N. Rouder and R. Ratcliff, “Comparing Exemplar- and Rule-Based Theories of Categorization,” *Current Directions in Psychological Science*, vol. 15, 2006.
- [24] S. Natal, I. McLaren, and E. Livesey, “Generalization of Feature- and Rule-based Learning in the Categorization of Dimensional Stimuli: Evidence for Dual Processes Under Cognitive Control,” *J Exp Psychol Anim Behav Process*, vol. 39, no. 2, pp. 140–51, 2013.
- [25] I. Gibbs, Bobby E., J. G. W. Bernstein, D. S. Brungart, and M. J. Goupell, “Effects of better-ear glimpsing, binaural unmasking, and spectral resolution on spatial release from masking in cochlear-implant users,” *The Journal of the Acoustical Society of America*, vol. 152, no. 2, pp. 1230–1246, 08 2022.