



Quantifying the effect of speech pathology on automatic and human speaker verification

Bence Mark Halpern^{1,2}, Thomas Tienkamp^{3,4}, Wen-Chin Huang¹, Lester Phillip Violeta¹, Teja Rebernik³, Sebastiaan de Visscher⁴, Max Witjes⁴, Martijn Wieling³, Defne Abur³, Tomoki Toda¹

¹Nagoya University, Japan ²Netherlands Cancer Institute, The Netherlands

³University of Groningen, The Netherlands ⁴University Medical Center Groningen, The Netherlands

halpern.bence.e8@f.mail.nagoya-u.ac.jp, tomoki@icts.nagoya-u.ac.jp

Abstract

This study investigates how surgical intervention for speech pathology (specifically, as a result of oral cancer surgery) impacts the performance of an automatic speaker verification (ASV) system. Using two recently collected Dutch datasets with parallel pre and post-surgery audio from the same speaker, NKI-OC-VC and SPOKE, we assess the extent to which speech pathology influences ASV performance, and whether objective/subjective measures of speech severity are correlated with the performance. Finally, we carry out a perceptual study to compare judgements of ASV and human listeners. Our findings reveal that pathological speech negatively affects ASV performance, and the severity of the speech is negatively correlated with the performance. There is a moderate agreement in perceptual and objective scores of speaker similarity and severity, however, we could not clearly establish in the perceptual study, whether the same phenomenon also exists in human perception.

Index Terms: speaker verification, speech pathology, oral cancer

1. Introduction

Automatic speaker verification (ASV) has seen a large development and is a crucial part of biometric systems. Recent studies, however, found that features used by ASV systems are useful for other novel applications. One example is predicting the severity of voice and articulation disorders [1, 2]. Given that ASV systems are designed to distinguish individuals based on unique voice features, it is not surprising that ASV learns some pathology-related features. However, this realisation strongly hints that the effectiveness of ASV systems might be reduced with pathological speakers, and raises important questions. For example, is the severity of a speech disorder negatively correlated with the performance?

Another interesting application case is in the perceptual experiments of voice conversion (VC) [3, 4]. In VC, an important aspect of performance evaluation is ensuring the converted voice's perceptual similarity to the target speaker's voice. This application case further raises the question: if the ASV system can be influenced by speech pathology, is human perception of speaker identity also influenced? If this phenomenon is not observable in human perception, ASV systems might not be an appropriate replacement for perceptual experiments in VC.

Despite these valid questions and the fact that there is a growing concern in the speech community about how models perform for atypical speech [5], the impact of speech pathology on ASV remains largely unexplored. Previous research on the effect of pathological ASV has provided valuable insights but also left critical questions unanswered. For instance, the work of [6] showed that patients with conditions affecting the larynx

face most challenges with ASV, in contrast to articulation, and phonation disorders. However [6] bases its analysis on cross-sectional data, which cannot exclude the difference between the speaker populations as a confound. While [7, 8] provided parallel data and reported similar results, it still remains unclear whether ASV speaker similarity scores align with those provided by human listeners.

Given these considerations, in this study, we evaluate the state-of-the-art (SotA) ECAPA-TDNN's [9] performance on parallel pathological and healthy¹ data using two recently collected Dutch datasets: NKI-OC-VC [8] and SPOKE [10]. Our aim is to provide a clearer understanding of how speech pathology affects ASV systems, identify their appropriateness for their current use cases (biometric, severity evaluation, perceptual), and based on our results, show potential pathways for improving their robustness and accuracy for these use cases.

Our research questions are as follows:

- RQ1** Is ASV performance (as measured by the equal error rate) reduced due to speech pathology (specifically, one caused by oral cancer surgery)?
- RQ2** If speech pathology affects ASV, is the severity of a speaker correlated with that speaker's ASV performance (as measured by the equal error rate)?
- RQ3** In general, what is the relationship between objective/subjective speaker similarity and objective/subjective severity difference of two utterances from the same speaker with different speech severity? Most importantly, if speech pathology affects ASV performance, is this phenomenon also observed in perceptual experiments?

2. Datasets

2.1. NKI-OC-VC

The NKI-OC-VC dataset [8] includes Dutch pathological speech from 6 oral cancer (OC) speakers (four male, two female) who had undergone a composite resection (COMANDO) surgery or comparable treatment for mostly advanced tongue tumours.

Data was collected from the participants at a maximum of three time points: before the surgery (T1), within a month after the surgery (T2), and approximately six months after surgery (T3). The recordings took place during scheduled speech therapy sessions. Participants were asked to read the Dutch text "Jorinde en Joringel" [11] consisting of 92 sentences during the recording session. The total duration of all speech recordings, across all speakers, was approximately 2.5 hours. One record-

¹We refer to this speech as 'healthy' in the rest of the study, but it is important to note that the tumour's presence can already impede speech production in pre-operative speech.

ing session (speaker/time point) lasted five minutes on average. In some cases, patients felt the experiment was difficult, in that case, we prematurely stopped the experiment.

The speech was recorded with a Roland R-09HR field recorder at 44.1 kHz sampling frequency and 24-bit depth. This was later downsampled to 16 kHz and quantized to 16-bit to match the requirement for the speaker embedding extractions. The dataset includes speech severity labels provided by Speech Language pathologists (SLPs) using a five-point Likert scale with 5 meaning healthy, and 1 meaning severe.

2.2. SPOKE

The SPOKE dataset collected by [10] was created to assess the acoustics and kinematics of individuals scheduled for surgery for early-stage lateral tongue tumours. The dataset contains Dutch speech data from a total of six participants (three male, three female). The speech was collected before the surgery (T1), and approximately six months after the surgery (T3). The speech was recorded in a sound-dampened booth [12]. Participants read 10 unique sentences during the recording session. The total duration of all speech recordings is approximately 20 minutes. One recording session (speaker/time point) lasted 1.5 minutes on average.

The speech was recorded using a Shure MX-153-T microphone with a sampling frequency of 22,050 Hz at 16-bit depth. After the recording, the audio was downsampled to 16 kHz and mixed to mono to match the requirements for the speaker embedding extraction.

3. Methods

3.1. Objective speaker similarity (ASV) sim_{obj}

For calculating the objective speaker similarity between two utterances, we extracted speaker embeddings using the ECAPA-TDNN [9] model from SpeechBrain [13].² In the trials, these embeddings were compared using the cosine similarity, which we term objective speaker similarity (sim_{obj}). All comparison trials were strictly on different utterances to avoid the impact of utterance similarity on the distributions.

3.2. Objective severity (P-ESTOI) sev_{obj}

For calculating the objective severity of the speech, we used P-ESTOI [14]. It involves creating a reference signal of intelligible speech from one or multiple healthy speaker(s), which is then time-aligned to the pathological speech signal using dynamic time warping (DTW). This alignment allows for a comparison of the signals and quantifies their divergence. P-ESTOI is demonstrated to yield high correlations with subjective intelligibility ratings, including OC speech [8]. For each utterance, we calculate the P-ESTOI score by either taking the corresponding T1 utterance as a reference (for T2 and T3 utterances) or taking the average of all T1 utterances (for T1 utterances).

3.3. Subjective similarity and severity

To obtain subjective scores, we carried out two subjective tests with 15 native Dutch raters on the SPOKE dataset. We could not carry out subjective speaker similarity experiments for the NKI-OC-VC due to restrictions on the data, and we used the subjective scores of SLPs. Participants had a mean age of 27.1

years (range = 24-30). For each test, we report the interrater correlation ($ICC_{2,k}$).

Subjective similarity (sim_{sub}) To obtain subjective similarity scores, we carried out an experiment akin to the four-point scale voice similarity experiment often done in voice conversion (VC) studies, where two audio samples are presented to the listener side-by-side [15]. However, for this study, we decided to adopt a 100 mm visual analogue scale (VAS) instead of the four-point scale that is common in VC studies. VAS is able to give a more nuanced range of uncertainty, which allows us to better understand speaker recognition phenomena.

It is important that the listeners were unaware of the fact that there could be pathological speech in the dataset, specifically that there could be healthy and pathological speech from the same speaker. The instructions they received were therefore as follows (translation from Dutch): “Please listen to the following two audio samples and rate them for speaker similarity. Please consider who is speaking according to the characteristics of the sound and then make a choice using the visual analogue scale that varies from “Different (0%)” to “Same (100%)” to rate the speaker similarity of the two audio samples. Please pick more extreme values if you are more confident in your choice.”

The stimuli for the test included 55 comparisons, 25 same speaker same time point, 25 same-speaker different time point comparisons, and 5 different-speaker same time point comparisons. The stimuli were sampled based on the objective scores to make sure that we tested a broad range of comparisons. We did not use the scores given to different speakers in the analysis, but these were still needed to provide an appropriate range for the listening test. The average of the obtained similarity scores (sim_{sub}) were used in our further experiments.

Subjective severity (sev_{sub}) Following this test, listeners re-listened to the same stimuli, now rating them for speech severity using the same VAS method. We used the following instructions, which have been shown to provide scores comparable to expert listeners [16]: “A “healthy” statement is 100%, by which we mean that the pronunciation is easy to understand, you could write down the utterance, if asked, without any difficulty, and the speaker has clear articulation with normal speaking rate. A severely “pathological” utterance is 0%, which means that the pronunciation is impossible to transcribe, the speaker has articulation problems, and the speaker may speak slower or faster than normal.”

4. Experiments

4.1. RQ1: Is automatic speaker verification impacted by the presence of pathology?

For RQ1, we assessed the impact of speech pathology on ASV performance. We examined the effect of pathology on same-speaker and different-speaker sim_{obj} distributions and measured ASV performance using the equal error rate (EER%).

To examine this effect, we break down sim_{obj} , based on the speaker’s identity and the timing of the recordings. For instance, comparisons between recordings from the same speaker at different times (e.g., T1 to T3) were labelled as “Same T1-T3”. We compared these with a baseline (“Same T1-T1” for same-speaker trials and “Diff T1-T1” for different-speaker trials), evaluating differences using a *Wilcoxon rank-sum test*.

For EER calculations, we grouped these trials to represent various levels of difficulty. The “T1-T1 EER” scenario, the simplest, involves comparisons within the Same T1-T1 and Diff

²<https://huggingface.co/speechbrain/spkrec-ecapa-voxceleb>

T1-T1 groups. The most challenging scenario, “Any EER”, includes all same-speaker (“Same Any”) and different-speaker (“Diff Any”) comparisons. Additionally, for the NKI-OC-VC dataset, we examined two specific EERs: T12 EER, excluding comparisons with T3, and T13 EER, excluding comparisons with T2. For the SPOKE dataset, T12 EER cannot be calculated, and T13 EER is the same as “Any EER”.

4.2. RQ2: Speaker-severity vs speaker EER

To investigate, whether severity is correlated with EER, we calculate “Any EER” for each speaker separately, and compare this with both the mean of objective severity scores sev_{obj} for each T3 speaker, and the mean of the subjective severity scores obtained in sev_{sub} .

4.3. RQ3: General relationship between variables

To understand the relationship between the subjective/objective speaker similarity and speech severity, we calculated a Pearson’s correlations matrix between all the following properties of all the utterance pairs that we rated in the SPOKE dataset: the subjective speaker similarity (sim_{sub}), the objective speaker similarity (sim_{obj}), the objective speech severity difference (Δsev_{obj}) and the subjective speech severity difference (Δsev_{sub}). The latter two is calculated as the absolute difference of the two scores.

5. Results

5.1. RQ1: Is ASV impacted by the presence of pathology?

Table 1: Equal error rate (%) when allowing for different time points in the same speaker comparison.

Dataset	T1-T1 EER (%)	T12 EER (%)	T13 EER (%)	Any EER (%)
NKI-OC-VC	1.00	6.35	3.50	5.83
SPOKE	0	-	-	0.13

Figure 1 shows the objective similarity distributions for both datasets. In the NKI-OC-VC, the distributions for same-speaker trials show higher scores compared to different-speaker trials, as expected. However, the presence of pathology shifts down the median of the distributions, most clearly in the Same T1-T2 comparisons, which suggests that the verification performance decreases most prominently at the T2 timepoint. All of the observed differences were significantly different from the T1-T1 reference distributions. Note, however, that the different-speaker distributions’ shift is marginal compared to the same-speaker distributions. The SPOKE dataset shown a similar trend, with same-speaker trials yielding higher cosine similarity scores than different-speaker trials. The impact of pathology is also evident in this dataset, as Same T1-T3 and Same Any have a broader distribution and lower median value. However, no statistical difference was observed between Same T1-T1 and Same T3-T3.

The results in Table 1 summarise the impact of speech pathology in terms of EER results. For the NKI-OC-VC dataset, the EER increases when we allow comparisons with pathological speech. The highest EER rates are observed in the T12 case. This result is probably caused by speech directly after the surgery being more affected than six months after surgery due to scarring.

Conversely, the SPOKE dataset shows an absence of error

rate (0%) when the comparison is within the pre-surgery time point (T1-T1 EER), but a slight increase is observed when post-surgery utterances are included in the comparison (Any EER), reflecting a subtle influence of pathology on the system’s performance. Again, the subtler increase can be accounted for the lower level of speech pathology in this dataset. We therefore conclude that speaker verification performance is impacted by speech pathology.

5.2. RQ2: Speaker-severity vs speaker-EER

The analysis (shown in Figure 2) revealed a significant negative correlation between the EERs and objective severity scores ($r = -0.87, p < 0.05$), indicating that as the estimated severity of the speech pathology increases, a higher EER result can be expected. Furthermore, the study found a moderate positive correlation between the EERs and the subjective severity ($r = 0.65, p < 0.05$). However, the correlation is (much) weaker compared to the objective scores, which might be due to the non-ideal scenario of combining SLP and non-expert listeners scores in the experiment. Therefore, we can conclude that the severity is at least moderately predictive of the EER performance for a given speaker.

5.3. RQ3: General relationship between variables

Table 2: Correlations between objective (sim_{obj}) and subjective (sim_{sub}) similarity measures, and objective (Δsev_{obj}) and subjective (Δsev_{sub}) severity difference measures. Asterisks indicate the significance level of the correlations: (*) $p < .05$, (**) $p < .01$, and (***) $p < .001$.

	sim_{obj}	sim_{sub}	Δsev_{obj}	Δsev_{sub}
sim_{obj}	1	-	-	-
sim_{sub}	0.54 (***)	1	-	-
Δsev_{obj}	-0.28 (*)	-0.15	1	-
Δsev_{sub}	-0.43 (**)	-0.16	0.40 (**)	1

In the subjective tests, we observed moderate inter-rater reliability with both the similarity ($r = 0.61$) and the severity ratings ($r = 0.58$).

The correlation analysis, summarised in Table 2 shows a moderate positive correlation between the objective speaker similarity measures and subjective similarity ratings ($r = 0.54$), suggesting that there is some level of agreement with the scores determined by the ASV and the human listeners.

The correlation between objective severity difference and subjective severity difference was not as high as expected ($r = 0.40$), which may suggest that severity is not a property that is consistently expressed or perceived at the utterance level. While P-ESTOI has been successful in evaluating the intelligibility of pathological speech at the speaker level [8, 14], its performance was less reliable at the utterance level [16], indicating that the features that P-ESTOI captures may not translate directly to the perceptual cues listeners use to judge severity on the utterance level.

The relation between the speech severity difference and the speaker similarity is somewhat complicated. First, there is a lack of significant correlation between the subjective severity difference and the subjective similarity ($r = -0.16$), and also the objective severity difference and the subjective similarity ($r = -0.15$). However, there still seems to be a moderate cor-

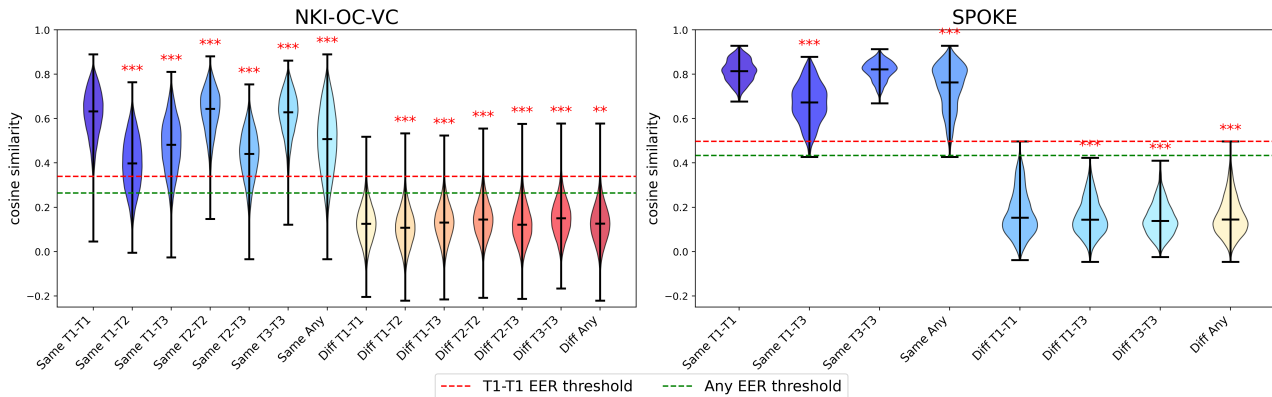


Figure 1: Comparison of cosine similarity distributions and EER thresholds for same-speaker and different-speaker trials across different time points. Asterisks indicate the significance level of the correlations: (**) $p < .01$, and (***) $p < .001$.

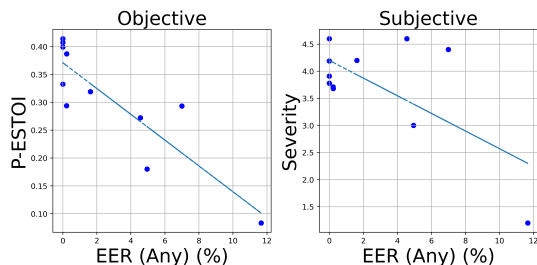


Figure 2: Objective and subjective speaker severity score in relation to EER performance.

relation between the subjective severity difference and objective similarity ($r = -0.43$), which complicates the conclusion here. This inconsistency can partly be due to rater bias, as well as the limited set of samples used in our study compared to previous research which reported a higher interrater correlation [16].

Overall, we conclude that there is a moderate agreement on the scores between the human and machine listeners but the relationship between subjective similarity and severity cannot be clearly established.

6. Discussion

We find that the ASV systems are affected by speech pathology, and performance deterioration is correlated to the severity of speech, independent of whether this severity is measured via objective or subjective measures. We found moderate agreement between perceptual and objective scores. However, we could not directly observe that speech severity affects subjective speaker similarity scores.

These findings have several implications. Firstly, current biometric systems do not seem to be robust to variations in speech pathology. Therefore, it is advised that future research focuses on training techniques that increase the robustness against these variations. Examples of successful strategies in mitigating these issues from other speech technology tasks point towards data augmentation using voice conversion [17] or other acoustic signal manipulations [18].

Secondly, it is clear from the experiment that current ASV models are also not well suited to replace perceptual experi-

ments. Humans will naturally exhibit different responses due to their different biases, therefore this is not surprising. A notable example of this is familiar speaker recognition, which is known to activate different pathways compared to non-familiar speaker recognition [19].

Nevertheless, replacing perceptual experiments for voice conversion is an important goal, and several strategies could be used to finetune the ASV system towards that goal. E.g. instead of speaker labels, directly using similarity scores from raters might be a better (but arguably, more cost-intensive) approach to train systems that model perceptual aspects well. A similar strategy has been already used and shown to improve the quality of synthetic speech [20]. Another way for development to consider is using listener-dependent modelling, which has been shown to be an effective strategy in speech naturalness ratings [21]. Listener-dependent modelling is a straightforward way to model this bias in the systems.

Thirdly, it is interesting that ASV scores are competitive with P-ESTOI, as the method outperformed the P-ESTOI on the utterance level. The EER reduction was also highly predictive of the severity of the speaker’s speech.

Finally, there are some limitations in our current analysis. First, our dataset comprises solely parallel data from low-to-mid speech severity speakers, where speaker embeddings are expected to be less impacted. Additionally, since the utterances used to compute the EER are from the same recording session, our estimation of the EER performance may be somewhat optimistic. Nonetheless, despite this potentially favourable scenario, we still observe a decline in performance, indicating that this limitation does not undermine the validity of our findings.

7. Conclusion

Our study showed that speech pathology negatively impacts automatic speaker verification (ASV) performance and performance reduction is negatively correlated with severity. We found moderate agreement between perceptual and objective scores, however, we could not establish the same phenomena in the perceptual experiments. The study highlights potential weaknesses in using these systems as biometrics and substituting perceptual experiments in voice conversion studies. We suggest strategies such as data augmentation, listener-dependent modelling, and similarity score training to alleviate these issues.

8. Acknowledgements

The data collection in the paper received ethical approval under the numbers IRBd20-159 (NKI-OC-VC), NL79242.042.21 (SPOKE), and ID95072353 (perceptual evaluation). The Department of Head and Neck Oncology and surgery of the Netherlands Cancer Institute receives a research grant from Atos Medical (Hörby, Sweden), which contributes to the existing infrastructure for quality of life research. This work was further supported by the Research School of Behavioral and Cognitive Neurosciences of the University of Groningen. This work is partly financed by the Dutch Research Council (NWO) under project number 019.232SG.011, and partly supported by JST CREST JPMJCR19A3, Japan.

9. References

- [1] S. Quintas, J. Mauclair, V. Woisard, and J. Pinquier, “Automatic Assessment of Speech Intelligibility using Consonant Similarity for Head and Neck Cancer,” in *Proc. Interspeech 2022*, 2022, pp. 3608–3612.
- [2] L. Moro-Velazquez, J. Villalba, and N. Dehak, “Using x-vectors to automatically detect parkinson’s disease from speech,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 1155–1159.
- [3] H. Wang, T. Thebaud, J. Villalba, M. Sydnor, B. Lammers, N. Dehak, and L. Moro-Velazquez, “Duta-vc: A duration-aware typical-to-atypical voice conversion approach with diffusion probabilistic model,” *arXiv preprint arXiv:2306.10588*, 2023.
- [4] H. Zhang, Z. Cai, X. Qin, and M. Li, “Sig-vc: A speaker information guided zero-shot voice conversion system for both human beings and machines,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6567–65 571.
- [5] S. Feng, O. Kudina, B. M. Halpern, and O. Scharenborg, “Quantifying bias in automatic speech recognition,” *arXiv preprint arXiv:2103.15122*, 2021.
- [6] S. T. Arasteh, T. Weise, M. Schuster, E. Noeth, A. Maier, and S. H. Yang, “The effect of speech pathology on automatic speaker verification: a large-scale study,” *Scientific Reports*, vol. 13, 2023.
- [7] J. Thienpondt, C. M. Speksnijder, and K. Demuynck, “Behavioral analysis of pathological speaker embeddings of patients during oncological treatment of oral cancer,” in *INTERSPEECH 2023*. ISCA, 2023, pp. 3018–3022.
- [8] B. M. Halpern, W.-C. Huang, L. P. Violeta, R. van Son, and T. Toda, “Improving severity preservation of healthy-to-pathological voice conversion with global style tokens,” in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2023, pp. 1–7.
- [9] B. Desplanques, J. Thienpondt, and K. Demuynck, “Ecapa-tdnn: Emphasized channel attention, propagation and aggregation in tdn based speaker verification,” 2020.
- [10] T. B. Tienkamp, T. Rebernik, R. Buurke, K. Polsterer, R. J. J. H. van Son, M. B. Wieling, M. J. H. Witjes, S. A. H. J. de Visscher, and D. Abur, “The effect of speaking style on the articulatory-acoustic vowel space in individuals with tongue cancer before and after surgical treatment,” in *Proceedings of the 13th International Seminar on Speech Production*, 2024, pp. 65–68.
- [11] R. J. J. H. van Son, D. Binnenpoorte, H. van den Heuvel, and L. C. W. Pols, “The IFA corpus: a phonemically segmented dutch “open source” speech database,” in *Proc. 7th European Conference on Speech Communication and Technology (Eurospeech 2001)*, 2001, pp. 2051–2054.
- [12] M. Wieling, T. Rebernik, and J. Jacobi, “SPRAAKLAB: a mobile laboratory for collecting speech production data,” in *Proceedings of the 20th International Congress of Phonetic Sciences (ICPhS 2023)*, 2023.
- [13] 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.
- [14] P. Janbakhshi, I. Kodrasi, and H. Bourlard, “Pathological speech intelligibility assessment based on the short-time objective intelligibility measure,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6405–6409.
- [15] Z. Yi, W.-C. Huang, X. Tian, J. Yamagishi, R. K. Das, T. Kinunen, Z.-H. Ling, and T. Toda, “Voice Conversion Challenge 2020 — Intra-lingual semi-parallel and cross-lingual voice conversion —,” in *Proc. Joint Workshop for the Blizzard Challenge and Voice Conversion Challenge 2020*, 2020, pp. 80–98.
- [16] B. M. Halpern, S. Feng, R. van Son, M. van den Brekel, and O. Scharenborg, “Automatic evaluation of spontaneous oral cancer speech using ratings from naive listeners,” *Speech Communication*, vol. 149, pp. 84–97, 2023.
- [17] Y. Zhang, Y. Zhang, B. M. Halpern, T. Patel, and O. Scharenborg, “Mitigating bias against non-native accents,” in *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, vol. 2022, 2022, pp. 3168–3172.
- [18] Z. Lin, T. Patel, and O. Scharenborg, “Improving whispered speech recognition performance using pseudo-whispered based data augmentation,” in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2023, pp. 1–8.
- [19] D. Van Lancker, J. Kreiman, and K. Emmorey, “Familiar voice recognition: Patterns and parameters part i: Recognition of backward voices,” *Journal of phonetics*, vol. 13, no. 1, pp. 19–38, 1985.
- [20] Y. Saito, S. Takamichi, and H. Saruwatari, “Perceptual-similarity-aware deep speaker representation learning for multi-speaker generative modeling,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1033–1048, 2021.
- [21] W.-C. Huang, E. Cooper, J. Yamagishi, and T. Toda, “Ldnet: Unified listener dependent modeling in mos prediction for synthetic speech,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 896–900.