

Noise-Robust Hearing Aid Voice Control

Iván López-Espejo, Eros Roselló, Amin Edraki, Naomi Harte, and Jesper Jensen

Abstract—Advancing the design of robust hearing aid (HA) voice control is crucial to increase the HA use rate among hard of hearing people as well as to improve HA users’ experience. In this work, we contribute towards this goal by, first, presenting a novel HA speech dataset consisting of noisy own voice captured by 2 behind-the-ear (BTE) and 1 in-ear-canal (IEC) microphones. Second, we provide baseline HA voice control results from the evaluation of light, state-of-the-art keyword spotting models utilizing different combinations of HA microphone signals. Experimental results show the benefits of exploiting bandwidth-limited bone-conducted speech (BCS) from the IEC microphone to achieve noise-robust HA voice control. Furthermore, results also demonstrate that voice control performance can be boosted by assisting BCS by the broader-bandwidth BTE microphone signals. Aiming at setting a baseline upon which the scientific community can continue to progress, the HA noisy speech dataset has been made publicly available.

Index Terms—Hearing aid, voice control, keyword spotting, noise robustness, bone-conducted speech.

I. INTRODUCTION

HEARING loss is a ubiquitous problem that, according to the World Health Organization, will affect 1 in 4 people by 2050 [1]. To mitigate the impacts of hearing loss, hearing aids (HAs) have the potential to help hard of hearing people improve their daily life engagement, confidence, and overall life quality [2]. Despite this, many hard of hearing people who could benefit from a HA do not wear one due to stigma and/or aesthetic reasons [3], [4]. Hence, advancing the design of invisible, in-ear-canal (IEC) HAs is key to increase their rate of use among people who need them.

To manipulate invisible HAs, voice control implemented via keyword spotting (KWS) [5] can be an extremely useful technology. In addition, voice control can also be a highly valuable tool to handle visible, behind-the-ear (BTE) HAs [6], [7], especially for elderly people who have a hard time manipulating the device’s small buttons [5].

Motivated by the above facts, this paper dives into essentially uncharted waters aiming at setting a baseline for noise-robust HA voice control upon which the scientific community can continue to advance. This is achieved through the following two contributions.

- 1) The creation of a one-of-a-kind speech dataset —based on the Google Speech Commands Dataset (GSCD) [8]— containing (simulated) noisy own voice as if it

were captured by a HA with 2 (front and rear) BTE microphones and 1 IEC microphone.

- 2) The training and evaluation of light, state-of-the-art KWS models [9] for HA voice control that leverage different combinations of microphone signals from the proposed HA noisy speech dataset.

In relation to the proposed dataset¹, which has been made publicly available², it is interesting to note that the IEC microphone is placed inside an ear occluded by an acoustic seal. This occlusion leads to bone-conduction vibrations originating from own voice to resonate inside the ear canal (*bone-conducted speech*, BCS). These vibrations are then captured by the IEC microphone with a relatively high signal-to-noise ratio (SNR), since BCS is not greatly contaminated by ambient noise due to the passive attenuation of the acoustic seal [11]. Nevertheless, due to transmission through human tissues, BCS suffers from a limited bandwidth (i.e., up to 2-2.5 kHz [10]) as well as from low-frequency distortion.

Given the noise-robustness properties inherent to BCS, assisting regular air-conducted speech by BCS is proven to boost the performance of several speech technology applications like speech enhancement [12], automatic speech recognition [13], and human sound classification [14]. In this work, we, for the first time, quantify the benefits of exploiting BCS for noise-robust HA voice control. Specifically, we demonstrate how a light KWS model ($\sim 50k$ parameters) [9] (which might be embedded into a HA) is able to provide a strong performance when employing the IEC microphone instead of the BTE ones, especially at highly-adverse SNRs. Furthermore, we also show that voice control performance can be significantly enhanced by complementing noise-robust BCS with the broader-bandwidth BTE microphone signals.

The rest of this paper is organized as follows. Section II explains how the HA noisy speech dataset is created. The KWS system used for HA voice control is described in Section III. KWS results on the HA speech dataset are discussed in Section IV. Finally, Section V concludes this work.

II. HEARING AID NOISY SPEECH DATASET

This section describes the proposed HA noisy speech dataset. First, the setup to measure HA-related transfer functions and how these are used to generate HA noisy own voice from mouth reference microphone (MRM) clean speech is presented in Subsecs. II-A and II-B, respectively. Then, Subsec. II-C explains how the dataset is created and structured.

While there are some similarities between HA datasets reported in [6], [7] and the one proposed here, this new corpus

This work has been funded by the Spanish Ministry of Science and Innovation under the “Ramón y Cajal” programme (RYC2022-036755-1).

I. López-Espejo and E. Roselló are with the Dept. of Signal Theory, Telematics and Communications, UGR, Spain (e-mails: {iloes,erosrosello}@ugr.es).

A. Edraki is with the Dept. of Electrical and Computer Engineering, QU, Canada (e-mail: a.edraki@queensu.ca).

N. Harte is with the School of Eng., TCD, Ireland (e-mail: nharte@tcd.ie).

J. Jensen is with Oticon A/S, Denmark, as well as with the Dept. of Electronic Systems, AAU, Denmark (e-mail: jesj@demant.com).

¹Note that the underlying hearing aid-related data originate from [10].

²<https://zenodo.org/records/14140800>




Fig. 1. Experimental setup for transfer function measurement depicting a subject wearing a left ear HA that embeds IEC, front and rear microphones (black, blue and red dots, respectively). The subject is surrounded by 16 equidistantly-spaced loudspeakers at eye-height (magenta dots).

is the only one to include both noisy speech and an IEC microphone.

A. Hearing Aid Transfer Function Measurement Setup

Own-voice transfer functions (OVTFs) and head-related transfer functions (HRTFs) were estimated from the experimental setup depicted in Fig. 1. One at a time, a total of 22 subjects sat in a soundproof sound studio while wearing HA prototypes³ by Demant A/S that embed 2 (front and rear) BTE microphones and 1 IEC microphone. Subjects were surrounded by 16 equidistantly-spaced loudspeakers laid out in a circular array at eye-height.

With this experimental setup, subjects read out random Wikipedia passages aloud so that an own-voice dataset could be collected [10]. Own voice was captured by all the HA microphones as well as by an MRM placed in front of each subject’s mouth. Then, subject-wise OVTFs between the MRM and each of the HA microphones were estimated using a least mean squares approach.

Additionally, sweep signals consisting of pure sinusoidal signals with frequency changing from near DC to slightly above 20 kHz were emitted by each loudspeaker around the subjects and recorded by the HA microphones. These signals were used to estimate subject-wise HRTFs between each loudspeaker and each of the HA microphones [15].

B. Noisy Own-voice Generation Model

Given a clean speech signal $x(n)$ captured at a HA user’s mouth, our aim is to simulate noisy versions of $x(n)$ at the IEC, front and rear HA microphones. We denote these noisy own-voice signals as $y_m(n)$, where $m \in \{i, f, r\}$, and i, f and r refer to the IEC, front and rear microphones, respectively.

Assuming a linear time-invariant system, the clean own-voice signal $x(n)$ at microphone m can be calculated, for all m , as $x_m(n) = h^m(n) * x(n)$, where $h^m(n)$ is a user-specific OVTf between the mouth and microphone m , and $*$ denotes

convolution. In addition, clean own voice is contaminated by additive noise coming from *any* L (which is an integer between 1 and 16) out of the 16 point noise sources represented by the loudspeakers of Fig. 1. Let $\{\nu_\ell(n)\}_{\ell=1}^L$ and $\{g_\ell^m(n)\}_{\ell=1}^L$ be the set of L noise signals at their sources and user-specific HRTFs between such sources and microphone m , respectively. Then, the noisy own-voice signal at microphone m , $y_m(n)$, is approximated, for all m , as

$$y_m(n) = x_m(n) + \alpha \sum_{\ell=1}^L g_\ell^m(n) * \nu_\ell(n), \quad (1)$$

where α is a noise scaling factor. Note that we do not account for the Lombard effect despite it is known to have an impact on speech recognition performance [16].

C. Dataset Creation and Structure

Mouth reference microphone signals $x(n)$ were obtained from the GSCD v2 [8], which is composed of more than 100k 1-second long utterances comprising one word each. Similar to [18], we first extracted clean training, validation and test sets out of the GSCD by picking all the signals in each set with an *a posteriori* SNR above 40 dB. Twelve classes (which are approximately balanced across the training, validation and test sets) were defined: the 10 standard keywords of the GSCD (“yes”, “no”, “up”, “down”, “left”, “right”, “on”, “off”, “stop” and “go”), the filler class (comprised of the other 25 GSCD words), and ambient noise (i.e., no own voice is present).

While we were provided with transfer functions from 22 subjects (see Subsec. II-A), only 5 of the subjects showed good ear occlusion (which was determined by visual inspection of HRTFs) and were kept for noisy own-voice generation. From those 5, 3 were allocated to the training set, 1 more, to the validation set, and the final one, to the test set. In turn, each GSCD training speaker was randomly assigned 1 out of the 3 training HA subjects.

Following the generation model of Subsec. II-B, the clean GSCD utterances were convolved with the corresponding OVTFs to create clean own-voice signals at the IEC, front and rear microphones. To improve the generalization ability of the neural models, *training* OVTFs were perturbed prior to convolution on an utterance- and microphone-basis as in [7]:

$$\tilde{h}^m(n) = (1 + \gamma_n)h^m(n) + \delta_n, \quad (2)$$

where $m \in \{i, f, r\}$, $\gamma_n \sim \mathcal{N}(\mu = 0, \sigma = 0.1)$ and $\delta_n \sim \mathcal{N}(\mu = 0, \sigma = 10^{-5})$.

Inspired by the structure of the classical AURORA-2 database [19], the training, validation and test sets were split into 3, 3 and 5 even partitions, respectively. Each of these partitions was distorted with a different noise type (see Table I) by considering Eq. (1), where *training* HRTFs were also perturbed as in Eq. (2). Copies of these partitions were generated at different front-microphone SNR levels by adjusting α in Eq. (1) according to the active speech level of $x_f(n)$ in line with the ITU P.56 recommendation [20]. Note that α is constant across $y_i(n)$, $y_f(n)$ and $y_r(n)$. Training and validation SNRs are $\{-15, -5, 5, 15, 25\}$ dB, whereas test SNR values are $\{-18, -9, 0, 9, 18\}$ dB. These SNR ranges

³For experimental purposes, and assuming symmetry, we just consider the left ear HA, as shown in Fig. 1.

TABLE I

DESCRIPTION OF THE TYPES OF NOISE USED TO GENERATE THE HA NOISY SPEECH DATASET. SSN STANDS FOR SPEECH-SHAPED NOISE. BABBLE, SSN, AND INTERFERING SPEAKER NOISES WERE BUILT FROM THE TIMIT DATASET [17] SUCH THAT TIMIT SPEAKERS DID NOT OVERLAP ACROSS THE TRAINING, VALIDATION AND TEST SETS OF THE PROPOSED HA DATASET

Noise type	Set	Description
Babble	Train, Val & Test	5 female and 5 male speakers randomly located (utterance-wise) around the HA user
Music	Train, Val & Test	Stereo songs played through two adjacent, randomly picked (utterance-wise) loudspeakers around the HA user
SSN	Train, Val & Test	Surround SSN played through all the loudspeakers around the HA user
Interfering speaker	Test	A female or male (same probability) speaker randomly located (utterance-wise) around the HA user
TV	Test	A TV program (America’s Got Talent) played through the loudspeaker at 0°




Fig. 2. Magnitude spectrograms (in dB) from an amplitude-normalized test utterance comprising the word “yes” which is contaminated by babble noise at a 0 dB SNR at the HA’s front microphone. From left to right: mouth clean reference, and IEC, front and rear microphones.

were roughly designed by taking into account expected own-voice and acoustic noise sound pressure levels as well as an approximate mouth-to-ear loss [21]. Finally, note that *every* utterance was contaminated with a distinct noise realization.

Roughly, the training, validation and test sets contain 18.8 hrs, 2.1 hrs and 2.4 hrs, respectively, of multi-channel audio.

For illustrative purposes, Fig. 2 displays magnitude spectrograms from a test utterance containing the word “yes” which is distorted by babble noise at a front-microphone SNR level of 0 dB. As can be seen, the IEC microphone signal comprising BCS is significantly less affected by noise but shows a more limited bandwidth than the front and rear microphone signals.

III. KEYWORD SPOTTING SYSTEM FOR VOICE CONTROL

Low-complexity KWS systems may be deployed in HAs to control them via voice [5]–[7]. In the next two subsections, the KWS model used in this work is introduced along with its training details.

A. Keyword Spotting Model

We use a light KWS model called BC-ResNet [9] that provides state-of-the-art results on the GSCD benchmark [5]. BC-ResNet is based on broadcasted residual learning which exploits the benefits of 1D temporal and 2D convolutions while assuring low computational complexity. In this paper, we test 6 variants of BC-ResNet which are denoted as BC-ResNet- τ , where $\tau \in \{1, 1.5, 2, 3, 6, 8\}$ refers to the factor by which the base number of model channels is multiplied. The larger τ is, the more parameters the model has.

BC-ResNet is fed with 40-dimensional log-Mel feature vectors computed by means of a 30 ms window with a 10 ms shift across 1-second long utterances. In case of evaluating HA microphone combinations, microphone-wise log-Mel feature matrices are stacked across the channel dimension before being input to the KWS model.




Fig. 3. Average KWS accuracy (%) as a function of the number of model parameters when testing the use of the different HA microphones. Shaded areas represent 95% confidence intervals.

B. Model Training Details

As in [9], BC-ResNet is trained for a total of 200 epochs to recognize the 12 classes stated in Subsec. II-C, i.e., 10 different keywords, filler and ambient noise. As an optimizer, stochastic gradient descent with a weight decay of 10^{-3} and a momentum factor of 0.9 is employed. During the first 5 epochs, the learning rate linearly increases from 0 to 0.1 as warm-up and then decays to 0 with cosine annealing [22]. The size of the mini-batch is 100 training utterances.

Finally, note that, different from [9], we consider neither SpecAugment [23] nor any other data augmentation method.

IV. EXPERIMENTS AND RESULTS

In order to establish a baseline for noise-robust HA voice control, this section presents KWS accuracy results from using the different HA microphones as well as their combinations.

TABLE II

AVERAGE KWS ACCURACY (%) WITH 95% CONFIDENCE INTERVALS AND RTF RESULTS OBTAINED BY BC-RESNET-3 EMPLOYING DIFFERENT HA MICROPHONES. KWS RESULTS ARE BROKEN DOWN BY SNR AND SEEN/UNSEEN NOISES. BEST RESULTS ARE MARKED IN BOLDFACE

Mic(s)	RTF	SNR (dB) – <i>Seen noises</i>					SNR (dB) – <i>Unseen noises</i>				
		-18	-9	0	9	18	-18	-9	0	9	18
IEC		50.7±2.6	86.0±0.6	92.9±1.0	94.5±0.4	95.1±0.4	37.2±2.6	81.3±0.9	90.6±0.3	93.6±0.8	94.2±0.4
Front	0.084	13.4±0.5	39.7±1.8	73.5±1.0	89.1±0.5	93.0±0.9	13.5±1.2	34.3±2.4	65.4±1.3	82.2±1.4	88.0±1.5
Rear		13.1±0.9	36.9±1.1	72.0±1.1	87.7±1.1	93.5±0.9	13.1±0.9	34.5±1.4	64.7±2.1	83.1±1.3	89.7±1.2
I+F	0.086	66.7±1.9	87.3±1.2	93.1±1.1	94.3±1.0	94.9±0.7	59.5±4.2	86.0±0.6	91.9±1.2	93.6±1.2	94.7±0.8
I+R		69.1±1.3	87.1±0.6	93.5±0.7	94.6±0.6	95.0±0.7	62.0±1.1	87.2±1.1	91.9±0.7	93.9±0.9	94.5±1.0
F+R		14.6±0.9	44.1±1.1	76.1±0.7	89.5±0.3	93.8±0.3	16.5±0.7	43.7±1.6	72.4±0.7	85.2±1.5	91.1±0.6
I+F+R	0.090	68.5±2.0	88.0±0.8	92.8±1.2	94.2±1.0	95.1±1.0	59.9±3.3	86.3±1.9	91.8±1.1	93.5±0.6	94.4±0.7




Fig. 4. Average KWS accuracy (%) as a function of the number of model parameters when testing HA microphone combinations. IEC curves are shown as a reference. Shaded areas represent 95% confidence intervals.

KWS accuracy (which is the standard metric for the GSCD [5]) is reported along with 95% confidence intervals from the Student’s t -distribution [24]. Given an experiment, to obtain its confidence interval, 5 distinct BC-ResNet models were trained by using different random parameter initialization.

Fig. 3 shows KWS accuracy (averaged across both SNRs and 5 trained models) as a function of the number of model parameters—which is determined by the value of τ (see Subsec. III-A)—when evaluating the use of the different HA microphones. The resulting curves are also broken down by whether noises were **seen** (babble, music, and SSN) or **unseen** (interfering speaker and TV) during training. As expected, KWS performance is better on the seen-noise test partition than on the unseen-noise one. Furthermore, while KWS accuracy resulting from using the front microphone is comparable to that from using the rear one, performance is clearly superior when employing the IEC microphone presumably due to the resilience of BCS to ambient noise.

Fig. 4 depicts KWS accuracy curves from the combination of the different HA microphones. IEC curves are also displayed as a reference. We can see from this figure that combining BCS from the IEC microphone with regular air-conducted speech from the BTE microphones boosts KWS performance. We conjecture that this is due to the noisy higher frequency information from the BTE microphone signals complementing noise-robust, bandwidth-limited BCS. Moreover,

we can also observe from Fig. 4 that I+F, I+R and I+F+R yield a comparable accuracy, which probably is due to the employed front- and rear-microphone spectral representations containing relatively similar information (see Fig. 2) given that the BTE microphones are close each other. This hypothesis is partially endorsed by F+R providing relatively small accuracy gains with respect to the front and rear microphones in isolation (see Fig. 4 versus Fig. 3).

Finally, note from Fig. 4 that KWS performance increases with the number of model parameters and starts to stagnate from 50k parameters. This makes BC-ResNet-3 (which has 54,168, 55,368 and 56,568 parameters for 1-, 2- and 3-microphone inputs, respectively) a good candidate to be deployed in HAs taking into account that they are low-resource devices. Table II shows detailed KWS accuracy results (broken down by SNR and seen/unseen noises) and real-time factors (RTFs)⁴ obtained by BC-ResNet-3 employing the previously tested microphones and combinations of them. From this table, we can clearly assess how the combination of noise-robust BCS with the broader-bandwidth BTE microphone signals yields good voice control performance at extremely challenging SNRs.

V. CONCLUSION

In this work, we have aimed at establishing a baseline for noise-robust HA voice control. This has been accomplished through 1) the creation of a publicly available² HA noisy dataset comprising both regular air- and bone-conducted own voice, and 2) the evaluation of light, state-of-the-art KWS models for HA voice control using such a dataset. Experimental results have demonstrated that it is possible to achieve robust voice control by exploiting BCS from an IEC microphone as well as voice control performance can be boosted by adding complementary noisy higher frequency information from BTE microphones.

Future work will investigate the collection of and evaluation on real noisy own voice, since the available own-voice data, which were used for OVTF estimation, are clean only.

⁴RTFs, defined as the ratio between the time taken to process an input and the input duration [25], were measured on an Intel Xeon Silver 4214R CPU with a clock frequency of 2.4 GHz.

REFERENCES

- [1] World Health Organization, “WHO: 1 in 4 people projected to have hearing problems by 2050,” 2021.
- [2] Ruth Jayné Avierinos, Eldre Beukes, Vinaya Manchaiah, Ilze Oosthuizen, Talita le Roux, and De Wet Swanepoel, “Meaningful life changes following hearing aid use: a qualitative user perspective,” *International Journal of Audiology*, pp. 1–10, 2024.
- [3] Macy Alcido and Maddie Lloyd, “Forbes Health Survey: Nearly Half Of People With Hearing Loss Believe There Is A Hearing Aid Stigma,” 2024.
- [4] Chloe Kent, “Hearing aid aesthetics: how the appearance of devices can impact adherence,” 2021.
- [5] Iván López-Espejo, Zheng-Hua Tan, John H. L. Hansen, and Jesper Jensen, “Deep Spoken Keyword Spotting: An Overview,” *IEEE Access*, vol. 10, pp. 4169–4199, 2022.
- [6] Iván López-Espejo, Zheng-Hua Tan, and Jesper Jensen, “Keyword spotting for hearing assistive devices robust to external speakers,” in *Proceedings of INTERSPEECH 2019 – 20th Annual Conference of the International Speech Communication Association, September 15-19, Graz, Austria*, 2019, pp. 3223–3227.
- [7] Iván López-Espejo, Zheng-Hua Tan, and Jesper Jensen, “Improved external speaker-robust keyword spotting for hearing assistive devices,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 1233–1247, 2020.
- [8] Pete Warden, “Speech Commands: A dataset for limited-vocabulary speech recognition,” *arXiv:1804.03209v1*, 2018.
- [9] Byeonggeun Kim, Simyung Chang, Jinkyu Lee, and Dooyong Sung, “Broadcasted residual learning for efficient keyword spotting,” in *Proceedings of INTERSPEECH 2021 – 22nd Annual Conference of the International Speech Communication Association, August 30-September 3, Brno, Czechia*, 2021, pp. 4538–4542.
- [10] Amin Edraki, Wai-Yip Chan, Jesper Jensen, and Daniel Fogerty, “Speaker adaptation for enhancement of bone-conducted speech,” in *Proceedings of ICASSP 2024 – 49th IEEE International Conference on Acoustics, Speech and Signal Processing, April 14-19, Seoul, Korea*, 2024, pp. 10456–10460.
- [11] Rachel E. Bouserhal, Tiago H. Falk, and Jérémie Voix, “In-ear microphone speech quality enhancement via adaptive filtering and artificial bandwidth extension,” *The Journal of the Acoustical Society of America*, vol. 141, pp. 1321–1331, 2017.
- [12] Ching-Hua Lee, Bhaskar D. Rao, and Harinath Garudadri, “Bone-conduction sensor assisted noise estimation for improved speech enhancement,” in *Proceedings of INTERSPEECH 2018 – 19th Annual Conference of the International Speech Communication Association, September 2-6, Hyderabad, India*, 2018, pp. 1180–1184.
- [13] Mou Wang, Junqi Chen, Xiao-Lei Zhang, and Susanto Rahardja, “End-to-end multi-modal speech recognition on an air and bone conducted speech corpus,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 513–524, 2023.
- [14] Liang Xu, Jing Wang, Lizhong Wang, Sijun Bi, Jianqian Zhang, and Qiuyue Ma, “Human sound classification based on feature fusion method with air and bone conducted signal,” in *Proceedings of INTERSPEECH 2022 – 23rd Annual Conference of the International Speech Communication Association, September 18-22, Incheon, Korea*, 2022, pp. 1506–1510.
- [15] Angelo Farina, “Simultaneous measurement of impulse response and distortion with a swept-sine technique,” in *The 108th Audio Engineering Society (AES) Convention, February 19-22, Paris, France*, 2000.
- [16] Ricard Marxer, Jon Barker, Najwa Alghamdi, and Steve Maddock, “The impact of the Lombard effect on audio and visual speech recognition systems,” *Speech Communication*, vol. 100, pp. 58–68, 2018.
- [17] J. S. Garofolo, L. F. Lamel, W. M. Fisher, J. G. Fiscus, D. S. Pallett, and N. L. Dahlgren, “DARPA TIMIT acoustic phonetic continuous speech corpus CDROM,” 1993.
- [18] Iván López-Espejo, Zheng-Hua Tan, and Jesper Jensen, “A novel loss function and training strategy for noise-robust keyword spotting,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 2254–2266, 2021.
- [19] David Pearce and Hans-Günter Hirsch, “The AURORA experimental framework for the performance evaluation of speech recognition systems under noisy conditions,” in *Proceedings of ICSLP 2000 – 6th International Conference on Spoken Language Processing, October 16-20, Beijing, China*, 2000.
- [20] ITU-T, “Objective measurement of active speech level,” Recommendation P.56, International Telecommunication Union, Geneva, Dec. 2011.
- [21] Dixita Patel, “Analyzing the Acoustics of a Head and Torso Simulator,” 2019.
- [22] Ilya Loshchilov and Frank Hutter, “SGDR: Stochastic Gradient Descent with Warm Restarts,” in *Proceedings of ICLR 2017 – 5th International Conference on Learning Representations, April 24-26, Toulon, France*, 2017.
- [23] Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk, and Quoc V. Le, “SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition,” in *Proceedings of INTERSPEECH 2019 – 20th Annual Conference of the International Speech Communication Association, September 15-19, Graz, Austria*, 2019, pp. 2613–2617.
- [24] N. Blachman and R. Machol, “Confidence intervals based on one or more observations,” *IEEE Transactions on Information Theory*, vol. 33, pp. 373–382, 1987.
- [25] Vineel Pratap, Qiantong Xu, Jacob Kahn, Gilad Avidov, Tatiana Likhomanenko, Awni Hannun, Vitaliy Liptchinsky, Gabriel Synnaeve, and Ronan Collobert, “Scaling Up Online Speech Recognition Using ConvNets,” in *Proceedings of INTERSPEECH 2020 – 21st Annual Conference of the International Speech Communication Association, October 25-29, Shanghai, China*, 2020, pp. 3376–3380.