

Multirate Audiometric Filter Bank for Hearing Aid Devices

Alice Sokolova^{*†}, Dhiman Sengupta[‡], Kuan-Lin Chen^{*}, Rajesh Gupta[‡], Baris Aksanli[†], Fredric Harris^{*}
Harinath Garudadri^{*}

^{*}Department of Electrical and Computer Engineering, UC San Diego, La Jolla, CA, USA

[‡]Department of Computer Science and Engineering, UC San Diego, La Jolla, CA, USA

[†]Department of Electrical and Computer Engineering, San Diego State University, San Diego, CA, USA

Email: {aasokolo, dhimnsen, kuc029, rgupta, fjharris, hgarudadri}@ucsd.edu; baksanli@sdsu.edu

Abstract—The frequency-dependent nature of hearing loss poses many challenges for hearing aid design. In order to compensate for a hearing aid user’s unique hearing loss pattern, an input signal often needs to be separated into frequency bands, or channels, through a process called sub-band decomposition. In this paper, we present a real-time filter bank for hearing aids. Our filter bank features 10 channels uniformly distributed on the logarithmic scale, located at the standard audiometric frequencies used for the characterization and fitting of hearing aids. We obtained filters with very narrow passbands in the lower frequencies by employing multi-rate signal processing. Our filter bank offers a 9.1× reduction in complexity as compared to conventional signal processing. We implemented our filter bank on Open Speech Platform, an open-source hearing aid, and confirmed real-time operation.

Index Terms—Filter bank, channelizer, hearing aids, multirate processing, pure tone audiometry

I. INTRODUCTION

Studies have shown that only about one-third of individuals who have hearing loss utilize a hearing aid. Among those individuals, around one-third do not use their hearing aids regularly. The main reason for this disuse is often the dissatisfaction with the speech quality offered by modern hearing aids, especially in noisy environments where hearing-impaired individuals need them the most [1]. Achieving music appreciation with hearing aids is an even greater challenge [2]. A wide variety of signal processing algorithms have been proposed to improve the sound quality of hearing aids, such as Speech Enhancement, Dynamic Range Compression, Frequency Warping, and more [3]–[6]. However, because human hearing is inherently frequency-dependent, many of these algorithms require dividing the input signal into frequency sub-bands [3]–[6], a process called sub-band decomposition.

Sub-band decomposition is prevalent in commercial and open-source hearing aid devices. Among open source hearing aids, one of the most widely employed designs is that of James

This work is supported by the National Institute of Health, NIH/NIDCD grants R21DC015046, R33DC015046, “Self-fitting of Amplification: Methodology and Candidacy,” and R01DC015436, “A Real-time, Open, Portable, Extensible Speech Lab,” awards to University of California, San Diego; grant IIS-1838830 from the Division of Information & Intelligent Systems, “A Framework for Optimizing Hearing Aids In Situ Based on Patient Feedback, Auditory Context, and Audiologist Input.”; San Diego State University “fred harris Excellence Scholarship; Qualcomm Institute at UCSD; Halicioğlu Data Science Institute at UCSD, and Wrethinking, the Foundation.”




Fig. 1. Frequency response of our proposed audiometric filter bank, shown on the logarithmic and linear scales.

Kates [7]. However, Kates’s baseline hearing aid algorithm, which contains six frequency bands, is starting to become insufficient for the advancing needs to audiometric research. Work in [8] determined that finer frequency subdivision can provide researchers with increased flexibility for hearing loss compensation, especially for unusual hearing loss patterns.

In this paper, we present a multirate real-time filter bank, also known as a channelizer, for audiometric research. Our proposed filter bank, shown in Figure 1, is uniformly distributed on the logarithmic scale, and reflects the behavior of the human cochlea. The cochlea inherently has frequency-dependent spectral resolution, meaning it can discern finer pitch variations at lower frequencies than at higher frequencies. Our proposed filter bank replicates this property by offering very narrow bandpass filters at lower frequencies, and wider filters at higher frequencies. Our channelizer also offers high sidelobe suppression and perfect signal reconstruction within ± 0.01 dB.

We implemented our filter bank on the Open Speech Platform (OSP) [9]–[11]. OSP is an open-source suite of software

and hardware tools for performing research in audiology, which includes a wearable hearing aid, a wireless interface, and a set of hearing enhancement algorithms.

II. OVERVIEW

Since the time of Gabor's work [12] in the middle of the 20th century, it is common knowledge among signal processing engineers that as the bandwidth of a filter decreases, the complexity of the filter increases. Reducing complexity is a crucial issue for hearing aids, which must operate in real-time and run off of battery power. We propose a filter bank which, to the best of our knowledge, offers higher frequency resolution than prior work while also offering real-time operation and low power consumption.

A. Center Frequencies

The structure of an audiometric filter bank reflects the nature of the human cochlea, which is inherently logarithmic. The American Speech-Language-Hearing Association (ASHA) defines a set of ten audiometric frequencies used for pure-tone audiometry, which are 8, 6, 4, 3, 2, 1.5, 1, 0.75, 0.5, and 0.25 kHz [13]. These are the frequencies most commonly targeted for audiometric filter banks.

These standard audiometric frequencies closely resemble a half-octave logarithmic sequence. However, every other frequency is not a true half-octave frequency, but rather, a simplified integer approximation. Our filter bank is a true half-octave channelizer, making it uniformly distributed on the logarithmic scale, as seen from Figure 1. Although some of the center frequencies don't exactly match the approximated standard ASHA frequencies, they are functionally the same, and for the sake of simplicity we will be referring to each individual band by its approximate audiometric frequency.

B. Attenuation and Ripple

The American National Standards Institute (ANSI) defines specifications for Half-Octave Acoustic filters [14]. The standard includes three classes of filters – class 0, 1, and 2, where class 0 has the tightest tolerances and class 2 has the lowest. Our filter bank meets class 0 standards – the highest of the three. Accordingly, each band of our filter bank has -75 dB sidelobe attenuation, and the in-band ripple is within ± 0.15 dB. The ripple of the composite response of our channelizer is also within ± 0.15 dB.

C. Filter Shape and Composite Response

One important property for an audiometric filter bank is perfect signal reconstruction. A filter bank has perfect reconstruction if the sum of all output bands is equal to the original input signal. In the frequency domain, perfect reconstruction implies that the composite frequency response of the filter bank (the sum of all magnitude responses) must be a flat line spanning all frequencies, as shown in Figure 7.

We ensure that our filter bank has perfect reconstruction by employing complementary filter design. Complementary filters are two filters whose sum is equal to an all-pass filter. For any

highpass or lowpass filter, its complement can be found by subtracting it from an all-pass filter. We employed complementary filters by first designing the rightmost highpass filter from Figure 1. We then found the complementary lowpass filter which forms the right edge of the second channel. We then designed another highpass filter for the left edge of the second channel, and convolved the two filters to form a single bandpass filter. This technique ensures that all adjacent channel edges are always complements of each other, which results in perfect signal reconstruction. Our channelizer offers perfect reconstruction within ± 0.01 dB.

Complementary filter design also resolves another design challenge. As seen from Figure 1b, a filter which is symmetrical on the logarithmic scale is asymmetrical on the linear scale, and visa versa. By constructing a bandpass filter using the convolution of a highpass and a lowpass filter, we were able to design asymmetrical filters which satisfy symmetry on the logarithmic scale.

III. MULTIRATE FILTERING

As mentioned earlier, the greatest challenge in designing a filter bank for hearing aids is reducing the complexity of the narrow bandpass filters at lower frequencies. We reduced the length of our narrow bandpass filters significantly by employing multirate signal processing.

Table I lists the number of taps needed to implement the filters shown in Figure 1 at a single sampling rate of 32 kHz. As the filters become narrower and sharper, they require an increasing number of taps, reaching impractical values at the lowest frequencies.

However, the complexity of a filter can be decreased by reducing the sampling rate. For any fixed frequency interval, the bandwidth of the interval is narrower relative to a higher sampling rate, and wider relative to a lower sampling rate. Thus, a filter spanning a fixed range of frequencies becomes relatively wider as the sampling rate decreases. As the relative filter bandwidth increases, the numbers of taps proportionately decreases. For example, when the sampling rate of a filter is decreased by half, the relative bandwidth of the filter doubles, and the number of taps needed to implement it is also halved.




Fig. 2. Dotted red lines represent the $1/2$, $1/4$, and $1/8$ sampling rates. Blue shading represents frequency bands which are processed at the original sampling rate. Green, yellow, and red correspond to $1/2$, $1/4$, and $1/8$ original sampling rate processing, respectively.




Fig. 3. Functional multirate filtering block diagram: The input signal is separated into multiple sampling rates, passes through the filter bank, and all outputs are restored back to the original sampling rate.

Due to the unique specifications of our audiometric channelizer, the filters at lower frequencies are narrower and sharper than the filters at higher frequencies. Thus, it is possible to map the ten filters in the channelizer to different sampling rates in order to reduce the complexity. Figure 2 shows the mapping of bands to sampling rates, where each color represents a sampling rate, and the dotted lines represent the Nyquist frequencies at each sampling rate. If a band lies to the left of a dotted line, it can be downsampled to the respective sampling rate, as long as sufficient transition bandwidth is reserved for an anti-aliasing filter. We will discuss anti-aliasing and resampling in more detail in Section III-A. Because of the logarithmic nature of our audiometric filter bank, it is convenient to map one octave per sampling rate and perform resampling in ratios of two.

Table I compares the single-rate versus multirate implementation of our channelizer. In the single-rate case, as the bandwidth of the filters is halved for every octave, the number of filter coefficients doubles for every octave. However, in the multirate implementation, we retain constant filter complexity because the decrease in a filter's bandwidth is compensated by a decrease in the sampling rate. Although the filters in our multirate filter bank can have different lengths, we adjusted all of them to be the same length for convenience.

Figure 3 shows the block diagram of our proposed Cascaded Multirate Filter Bank. The incoming frame enters at 32 kHz and gets separated into four different sampling rates using downsamplers. The multirate input frames then pass through the bandpass filters. Then each output is converted back to the original sampling rate using upsamplers. The final stage of the

TABLE I
NUMBER OF FILTER TAPS PER BAND WITH AND WITHOUT MULTIRATE PROCESSING

Filter Band:	Filter Taps		Sampling rate
	Single-rate	Multi-rate	
8 kHz	93	93	1
6 kHz	93	93	1
4 kHz	186	93	1/2
3 kHz	186	93	1/2
2 kHz	372	93	1/4
1.5 kHz	372	93	1/4
1 kHz	744	93	1/8
0.75 kHz	744	93	1/8
0.5 kHz	744	93	1/8
0.25 kHz	744	93	1/8

Multirate Filter Bank is a delay within each band, which will be discussed in more detail in Section III-C. The cascaded downsampling structure benefits from computational reuse.

A. Resamplers

In the previous section, we discussed the benefits of multirate filtering. However, multirate signal processing comes at the cost of resamplers which, if not designed efficiently, will significantly diminish the advantages of multirate processing.

As seen from Figure 3, our Multirate Filter Bank contains downsamplers with a ratio of 2:1, and upsamplers with a ratio as high as 1:8. Conventionally, downsampling is performed by passing a signal through an antialiasing filter and then discarding unnecessary samples, and upsampling is performed by zero-padding a signal and then passing it through an interpolating filter. As such, the complexity of conventional resamplers strongly depends on their resampling ratio – a high-ratio downsampler would require a sharp anti-aliasing filter to remove all unwanted frequencies, and a high-ratio upsampler would require a sharp interpolating filter to remove spectral signal copies. We avoid including sharp anti-aliasing and interpolating filters by cascading multiple resamplers of ratio 2 to achieve the desired resampling ratio.

Figure 4 compares a single-stage and a cascaded implementation of a 1:8 upsampler. A 1/8 band filter suitable for this resampler would require about 261 taps. The number of




Fig. 4. For a 1:8 upsampler, it is more efficient to use a cascade of three 1 : 2 resampling filters than a single stage filter.

Conventional vs. Polyphase Downsampler




Fig. 5. A workload comparison between a conventional and an equivalent polyphase implementation of a 2:1 downsampler. (All even coefficients of a halfband filter are zero except the middle coefficient.)

multiply-and-add operations, equal to the frame size multiplied by the number of filter coefficients, would equal to 8352 operations per 32-sample output frame. However, this upsampler can be split into three 1:2 upsamplers, each containing a half-band filter. Moreover, after each upsampling stage, the transition bandwidth of the interpolating filter can be increased, which reduces its complexity. Multiplying frame size by filter length yields only 680 multiply-and-add operations for a cascaded implementation.

The power consumption of the resamplers can be further reduced by utilizing polyphase filters [15]. Polyphase filters split the resampler into multiple paths and employ the Noble Identity such that the anti-aliasing or interpolating filter always runs at the lower sampling rate. Figure 5 compares the conventional and polyphase implementations of a 2:1 downsampler. The workload of the conventional filter is 1120 multiply-and-add operations per 32-sample input frame. Polyphase filtering reduces the complexity by rearranging filtering and down-sampling. Moreover, by definition, each even coefficient of a halfband filter is equal to zero except the middle coefficient. As such, the top branch of the polyphase filter in Figure 5 reduces to only one coefficient, and the total number of multiply-and-add operations is reduced from 1120 to 304 per 32-sample frame.

B. Power

We estimate the cumulative power consumption of our filter bank by computing the total number of multiply-and-add operations per one output sample. We compute the number of operations per sample of our multirate channelizer by calculating the total number of operations per a 32-sample input frame, and then normalizing by frame size. The number of operations at each stage is found by multiplying current frame size by filter length. Due to the multirate structure of our channelizer, normalization by frame size may result in a fractional number of operations per sample.

Table II compares the total number of multiply-and-accumulate operations per sample for a single-rate and multirate implementation of our channelizer. The multirate operations estimate accounts for all filters and resamplers. Our evaluations show that compared to a conventional approach, our

TABLE II
TOTAL MULTIPLY-AND-ADD OPERATIONS PER SAMPLE PER BAND (TOTAL OPERATIONS PER FRAME NORMALIZED BY FRAME SIZE)

Filter Band:	Operations Per Sample		Ratio
	Single-rate	Multi-rate	
8 kHz	93	93	1x
6 kHz	93	93	1x
4 kHz	186	65.5	2.8x
3 kHz	186	56	3.3x
2 kHz	372	41	9.1x
1.5 kHz	372	31.5	11.8x
1 kHz	744	34.9	21.3x
0.75 kHz	744	18.3	40.8x
0.5 kHz	744	18.3	40.8x
0.25 kHz	744	18.3	40.8x
Total	4278	469.6	9.1x

multirate filter bank offers 9.1 $\times$ improvement in complexity. For a wearable battery-operated system, power consumption and processing capabilities are of critical importance. Reducing the number of operations improves battery-life and frees processing power for other tasks.

C. Delay

As seen from Figure 3, different frequency bands follow different signal paths and as such, experience varying amounts of delay. Because of the resamplers and lower sampling rates, lower frequency bands incur more delay than higher frequencies. Inserting a delay at the end of each output band, as seen in Figure 3, gives us the the option of aligning the bands and ensuring uniform group delay.

Figure 6a shows the unaligned impulse responses of our channelizer. The highest frequency channels experience about 1.5 ms of delay, and the 3 kHz and 4 kHz bands, where the majority of speech content resides, experience about 4 ms of delay. However, the latency disparity causes a phase offset between the ten bands, which leads to distortion in the composite frequency response, as seen in Figure 7a. To certain listeners, this phase disparity sounds like an echo or a distorted sound timbre.

This phase disparity can be resolved by delaying the higher frequency bands until the lower frequency bands arrive. Figure 6b shows the aligned impulse responses of our filter bank. Aligning the bands eliminates the distortion in the composite frequency response and recovers perfect reconstruction, as seen from Figure 7b.

Conventionally, the latency limit for a real-time hearing aid is considered to be around 10 milliseconds [16]. As seen from Figure 6b, the latency of the aligned channelizer is about 18 ms, which exceeds the conventional real-time threshold. However, as of the writing of this paper, we are actively developing a minimum phase implementation of our filter bank, which offers all the features of a linear phase implementation, including perfect reconstruction, but with much smaller latency. Our preliminary experiments are showing promising results, with an overall latency of less than 5 ms.




Fig. 6. (a) Shows the unaligned impulse responses of the ten-band filter bank, and (b) shows the aligned impulse responses. As seen in (a), the higher frequency bands have less delay than the lower frequency bands. The signals are aligned by delaying the higher frequency bands.

IV. SYSTEM INTEGRATION

We have integrated our proposed filter bank into the Open Speech Platform (OSP) developed at University of California San Diego [10], [11], [17]. OSP is an open source suite of hardware and software tools for conducting research into many aspects of hearing loss both in the lab and the field. The hardware system consists of a battery operated wearable device running a Qualcomm 410c processor, similar to those in cellphones, with two ear-level assemblies attached – one for each ear. More details about the hardware systems can be found in [10].

At the core of OSP software is the real-time Master Hearing Aid (RT-MHA) reference design. Initially, the incoming audio signal from the microphones is sampled at 48 kHz, and is then downsampled to 32 kHz (not to be confused with the resamplers present in the channelizer). The audio signal is then routed to the channelizer. The outputs of the channelizer then pass through the Wide Dynamic Range Compression (WDRC) unit to compensate for the user's hearing loss. Then the outputs of the WDRC are recombined and passed through a Global Maximum Power Output (MPO) controller in order to limit the power outputted by the speaker. Finally, the audio is upsampled from 32 kHz back to 48 kHz and outputted through the speakers. Additionally, the RT-MHA reference design contains Adaptive Feedback Cancellation (AFC) in order to compensate for the feedback arising from the close proximity of the microphone and the speaker. More detailed explanations of the RT-MHA components can be found in [10], [11].

We evaluated the performance of our proposed filter bank by running three different MHA designs on the hardware system for 15 minutes: the 6-band reference design, the 10-band unaligned design, and the 10-band aligned design. For each design, we measured the amount of time it took to process 1 millisecond of audio data. As long as all the processing is




Fig. 7. (a) Shows the frequency response of the unaligned composite filter, and (b) shows the frequency response of the aligned composite filter. As expected, the aligned composite filter has a flat spectrum.

completed within 1 millisecond before the next frame of data arrives, the system can operate in real-time.

Table III shows the mean and standard deviation of the processing times of each component for all three MHA designs. Our experiments show that OSP running with our proposed filter bank meets the real-time processing deadline of 1 millisecond. Moreover, the 10-band aligned and unaligned channelizers outperform the 6-band channelizer by about 65 microseconds, offering more frequency bands for less processing time. Most other stages of the MHA reference design were unaffected by the transition from 6 bands to 10 bands, with the exception of WDRC. However, despite the increased workload for the WDRC unit due to an increased number of channels, the overall processing time for the 10-band design is statistically similar to that of the 6-band design.

V. COMPARISON WITH PRIOR WORK

We compared our proposed audiometric filter bank against the filter bank from the widely popular Kates digital hearing

TABLE III
REAL-TIME PROCESSING PERFORMANCE STATISTICS FOR THE 6 BAND FILTER BANK IMPLEMENTATION [10] AND 10 BAND IMPLEMENTATION PRESENTED IN THIS PAPER. (TOTAL TIME TAKEN TO PERFORM EACH PROCESSING STEP ON A 1 MS AUDIO BUFFER.)

	6 Band [10]		10 Band Unaligned (μs)		10 Band Aligned (μs)	
	Mean	STD	Mean	STD	Mean	STD
Down Sampling	14.77	1.23	14.83	1.33	14.64	1.11
Filter Bank	175.51	3.49	104.365	3.89	108.55	4.55
WDRC	60.12	2.09	100.42	3.23	100.80	3.21
Global MPO	11.48	0.71	11.47	0.69	11.49	0.80
AFC	131.72	2.78	131.96	2.81	131.98	2.97
Up Sampling	14.57	1.03	14.67	1.14	14.73	1.21
Overall HA*	510.82	64.91	490.68	65.37	487.81	64.92

*The measured total time includes overheads like audio processing callback, including sending work to the threads for the left and right channels and waiting for them to complete.




Fig. 8. Frequency responses of our proposed audiometric filter bank and the Kates six-band filter bank.

aid [7]. Kates' hearing aid is the common go-to system for hearing aid research. Our proposed filter bank meets and improves upon the capabilities of the Kates system.

Figure 8 presents our proposed channelizer and the Kates filter bank on a logarithmic scale. Our proposed filter bank offers four more bands than Kates' filter bank. These additional bands provide our channelizer with higher spectral resolution, which makes it possible to more accurately fit a hearing aid to a particular hearing loss prescription.

Within each band, our filter bank also offers higher band integrity, especially at lower frequencies. The lower frequency bands of Kates's filter bank have very wide side-lobes, which introduces cross-talk between bands. According to the ANSI Standard on Half-Octave Filters [14], Kates's 500 Hz band does not even fall within class 2 standards, whereas all the filter in our proposed channelizer are class 0.

Across-band cross-talk plays a detrimental role in the ability of a hearing aids algorithm in faithfully fulfilling a hearing aid prescription. Based on a user's hearing loss pattern, determined using pure-tone audiometry [13], a hearing aid prescription determines the desired output levels for various signal input levels at different audiometric frequencies. When cross-talk from neighboring bands interferes with any given band, the hearing aid will no longer match the desired output levels outlined in the prescription.

We compared the accuracy of prescription matching of our filter bank with the Kates filter bank by running the OSP master hearing aid (MHA) using both filter banks. Both the MHA algorithms were calibrated to take into account the frequency response of the physical microphone and speaker. Each of the master hearing aids were loaded with an identical prescription. Then using a hearing aid verification device called the Verifit2 [18], we measured the gain at five frequencies over ten different input levels. This allowed us to verify how faithfully each MHA algorithm was able to reproduce the prescription.




Fig. 9. Accuracy of hearing loss prescription matching of our proposed filter bank vs. Kates six-band filter bank.

TABLE IV
COMPUTATIONAL WORKLOAD COMPARISON BETWEEN PROPOSED 10-BAND FILTER BANK AND KATES 6-BAND FILTER BANK

Filter Bank:	Bands:	Operations per sample
Proposed OSP Filter Bank	10	469.6
Kates Filter Bank	6	1542

Figure 9 shows the error from the ideal hearing aid prescription as measured by the Verifit2. Error within 3 dB of the target prescription is considered acceptable by the audiologist community because according to [19] any error lower than 3 dB is not noticeable by the user, though lower error is always desirable. As seen from Figure 9, both filter banks perform well, but the Kates filter bank sometimes crosses the 3 dB threshold at lower frequencies. Comparatively, our proposed filter bank has lower error overall, especially at lower frequencies, and never crosses the 3 dB threshold.

Our proposed filter bank, despite having more bands than the Kates filter bank, also has lower complexity due to multirate processing. Table IV compares the operations per sample of our proposed filter bank and the Kates filter bank. Every filter in the Kates system contains 257 coefficients, bringing the total to 1542 operations per sample for six bands. Comparatively, our proposed filter bank requires 469.6 operations per sample, as shown in Table II – more than 3× less than the Kates system.

VI. CONCLUSION

In this paper, we presented a 10-band multirate filter bank for audiology and hearing aids research. Our filter bank is tailored to the needs of the audiology community, which requires a channelizer that reflects the nature of the human cochlea, which has frequency-dependent spectral resolution. Our proposed filter bank targets the audiometric frequencies used for characterizing hearing loss in pure tone audiome-

try [13]. The individual bandpass filters in our channelizer offer at least 75 dB of sidelobe attenuation, and the composite response offers near perfect signal reconstruction with no more than 0.01 dB ripple in the pass-band. In order to lower the power consumption of the very narrow filters present in our channelizer, we employed multirate signal processing, which resulted in a $9.1\times$ reduction in complexity compared to single-rate processing. We integrated our channelizer into Open Speech Platform [9], and verified that the Master Hearing Aid algorithm running our filter bank meets the real-time processing deadline and outperforms the previous channelizer used in the Open Speech Platform.

REFERENCES

- [1] R. J. Bennett, A. Laplante-Lévesque, C. J. Meyer, and R. H. Eikelboom, "Exploring hearing aid problems: Perspectives of hearing aid owners and clinicians," *Ear and Hearing*, vol. 39, no. 1, pp. 172–187, 2018.
- [2] M. Chasin and F. A. Russo, "Hearing aids and music," *Trends in Amplification*, vol. 8, no. 2, pp. 35–47, 2004.
- [3] Y. Hu and P. C. Loizou, "Subjective comparison and evaluation of speech enhancement algorithms," *Speech communication*, vol. 49, no. 7–8, pp. 588–601, 2007.
- [4] Y. Ghanbari, M. Karami, and B. Amelifard, "Improved multi-band spectral subtraction method for speech enhancement," in *Proc. 6th IASTED internat. conf. on signal image process*, pp. 225–230, 2004.
- [5] J. M. Kates, "Dynamic range compression using digital frequency warping," Sept. 6 2011. US Patent 8,014,549.
- [6] L. Lee and R. Rose, "A frequency warping approach to speaker normalization," *IEEE Transactions on speech and audio processing*, vol. 6, no. 1, pp. 49–60, 1998.
- [7] J. M. Kates, *Digital hearing aids*. Plural publishing, 2008.
- [8] P. E. Souza, "Effects of compression on speech acoustics, intelligibility, and sound quality," *Trends in amplification*, vol. 6, no. 4, pp. 131–165, 2002.
- [9] "THE lab at UC San Diego," *Open Speech Platform*, 2019. Available: <http://openspeechplatform.ucsd.edu/>.
- [10] L. Pisha, J. Warchall, T. Zubatiy, S. Hamilton, C.-H. Lee, G. Chockalingam, P. P. Mercier, R. Gupta, B. D. Rao, and H. Garudadri, "A wearable, extensible, open-source platform for hearing healthcare research," *IEEE Access*, vol. 7, pp. 162083–162101, 2019.
- [11] D. Sengupta, T. Zubatiy, S. K. Hamilton, A. Boothroyd, C. Yalcin, D. Hong, R. Gupta, and H. Garudadri, "Open speech platform: Democratizing hearing aid research," in *Proceedings of the 14th EAI International Conference on Pervasive Computing Technologies for Healthcare*, 2020.
- [12] D. Gabor, "Theory of communication. part I: The analysis of information," *Journal of the Institution of Electrical Engineers-Part III: Radio and Communication Engineering*, vol. 93, no. 26, pp. 429–441, 1946.
- [13] "Guidelines for manual pure-tone threshold audiometry," *American Speech-Language-Hearing Association*, Retrieved May 2020.
- [14] *Specification for Octave-Band and Fractional-Octave-Band Analog and Digital Filters*. ANSI Standard S1.11-2004.
- [15] R. E. Crochiere and L. R. Rabiner, *Multirate digital signal processing*. Prentice-Hall, 1983.
- [16] 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.
- [17] H. Garudadri, A. Boothroyd, C.-H. Lee, S. Gadiyaram, J. Bell, D. Sengupta, S. Hamilton, K. C. Vastare, R. Gupta, and B. D. Rao, "A realtime, open-source speech-processing platform for research in hearing loss compensation," in *2017 51st Asilomar Conference on Signals, Systems, and Computers*, pp. 1900–1904, IEEE, 2017.
- [18] "Audioscan Verifit 2," *Audioscan/Etymonic Design Inc.*, 2014. Available: <https://www.audioscan.com/en/verifit2/>.
- [19] D. McShefferty, W. M. Whitmer, and M. A. Akeroyd, "The just-noticeable difference in speech-to-noise ratio," *Trends in hearing*, vol. 19, p. 2331216515572316, 2015.