



Low-Level Physiological Implications of End-to-End Learning of Speech Recognition

Louise Coppieters de Gibson^{1,2}, Philip N. Garner¹

¹Idiap Research Institute, Martigny, Switzerland,

²Ecole polytechnique fédérale de Lausanne (EPFL), Switzerland

louise.coppieters@idiap.ch, pgarner@idiap.ch

Abstract

Current speech recognition architectures perform very well from the point of view of machine learning, hence user interaction. This suggests that they are emulating the human biological system well. We investigate whether the inference can be inverted to provide insights into that biological system; in particular the hearing mechanism. Using SincNet, we confirm that end-to-end systems do learn well known filterbank structures. However, we also show that wider band-width filters are important in the learned structure. Whilst some benefits can be gained by initialising both narrow and wide-band filters, physiological constraints suggest that such filters arise in mid-brain rather than the cochlea. We show that standard machine learning architectures must be modified to allow this process to be emulated neurally.

Index Terms: speech recognition, cochlear models, end-to-end architectures, filterbanks, SincNet

1. Introduction

Advances in automatic speech recognition (ASR) have led to performance that is very good in terms of word error rate (WER), but perhaps at the expense of our own understanding of how they function. End-to-end (E2E) techniques [1] have removed the need for knowledge of the hearing mechanism. Self-supervised training [2] has done the same for phonetics. More generally, large pre-trained models are available [3] removing the need for even the machine learning (ML) know-how. Given that these systems work well, the question arises: “what have they learned?” This is difficult to answer because their component parts cannot readily be mapped to biological ones.

In this study, we are interested in the hearing mechanism. The biological mechanism is quite well understood [4], with important parts being the logarithmic response to both frequency and amplitude. For many years, filterbank approaches were used as models of this process [5, 6]. Whilst many variations have been studied, the authors’ ad-hoc experience suggests that the details do not lead to big changes. Recent E2E approaches, however, have clearly demonstrated that training the filterbanks can be beneficial [7]. A (1D) convolution layer in the machine learning field is a filter in the signal processing field. However, the only validations of which we are aware tend to show that the component convolutions learn filters with a distribution similar to a mel filterbank. This in turn tends to reinforce the above question rather than answer it.

In SincNet, Ravanelli et al. [8] constrained the convolutions to be a sinc ($\sin(x)/x$) form, leading to a rectangular band-pass filter. The filter is then defined by two trainable parameters: the lower frequency and the bandwidth. Whilst not being biologically accurate, this approach has a distinct attraction of being interpretable.

In the remainder of this mainly experimental paper, we describe SincNet and a modest frame-based experimental scenario. We confirm that SincNet learns a mel filterbank, but also show that wider bandwidth filters are important for performance. We argue that such filters arise because of restrictions of standard ML convolutional architectures, and conclude with what this infers about how to construct a biologically plausible hearing model.

2. Background

The study of the human cochlea has interested many researchers since the beginning of the 20th century. Von Békésy laid the groundwork of the research on this topic in 1960 [9].

The basilar membrane in the cochlea can be interpreted as a natural filterbank [10, 11]. Current understanding of the working of the cochlea is that wave propagation is an active process [12] and works as an array of Hopf oscillators [13, 14]. However, in this study, we limit ourselves to passive analogues. The scaling of this natural filterbank has been analysed from different points of view, leading to several scaling (or warping, spacing) approaches. The Greenwood scaling [15] is the one that best represents the scaling of the frequency sensitive hair cells in the cochlea based on the physical distance on the basilar membrane of the hair cells. The mel scale [16] is based on frequencies judged to be equally spaced in human perceptual tests. Bark [17, 18] and ERB (equivalent rectangular bandwidth) [19] are somewhere between mel and Greenwood, but by contrast are derived from critical *bands* of hearing.

ASR frontends take either some preprocessed features as input or, more recently, raw input waveforms. Filterbanks have been the basis of feature extraction [20] for many years. As early as 2001, a study [21] showed that a filterbank could be obtained from a mathematical derivation of a data driven design. From the resulting filterbank, about half of the filters were then kept to represent the filterbank motivated by the fact that those filters were enough to cover the whole frequency range. For the E2E approaches, CNNs for ASR were introduced by Hinton et al. [22, 23] and have been used for a decade. Since 2018 some architectures propose a way to combine both the filterbanks and the E2E architecture, where the filterbanks are trainable and part of the convolution layers. Zeghidour et al. [24] proposed an implementation with Gabor filters and Ravanelli et al. [25] with rectangular filters (SincNet). Other work on trainable filterbanks includes that of Seki et al. [26], who proposed an architecture based on a filter layer combined with a DNN where the filter features were directly computed with a log-compression after the filter layer. In that study the gain, central frequency, bandwidth and filter shape were free to train, whilst in SincNet only two parameters are free to train, defining the filters in the first layer.

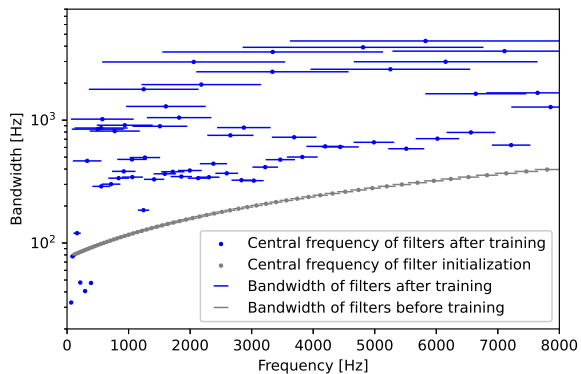


Figure 1: Evolution of the baseline implementation of SincNet: the graph on the top shows the initial filter distribution and the bottom plot show the filter distribution after training. The x-axis represents the frequency range and the y-axis the amplitude of the filters. The filters themselves are represented by their central frequency (dot) and their bandwidth (bar).

3. Initial Analysis

In an initial, quite basic analysis, the main motivation was to understand what the trainable filters learn; i.e., which typical hyperparameters (e.g., the number of filters needed to describe the signal) can be derived from those trainable filter models, which initializations are appropriate. For this study, we focus on SincNet.

3.1. SincNet setup

The SincNet model [25] is built with a 4-layer CNN followed by a 5-layer DNN. The first layer of the CNN is constructed with trainable filters. Those filters are initialised as a rectangular bandwidth mel-scale filterbank, an easily computable type of filter in the time domain. Since the inverse Fourier transform of a rectangular low-pass filter is a sinc function, a rectangular bandwidth filter can be written as the difference of two low-pass filters as in equation 1.

$$h[n] = \text{sinc}(2\pi f_2 n) - \text{sinc}(2\pi f_1 n), \quad (1)$$

where $h[n]$ represents a typical filter of the first convolutional layer. The trainable parameters of the SincNet filters are the lower frequency (f_1) and the bandwidth ($f_2 - f_1$), i.e., linear combinations of f_1 and f_2 . Moreover, in the time domain, filtering a signal is mathematically equivalent to the convolution of this signal with the filter kernel. Between the different convolution layers the following operations are used: maxpooling for downsampling, layernorm, ReLU and dropout before passing through a five-layer DNN. The input of signal is a raw waveform of 200 ms at 16 kHz. For this research the experiments are performed on TIMIT [27] and to verify that the observations are not database related, the baseline experiments have been double checked with the mini-Librispeech database [28].

3.2. Baseline

In the default implementation, the number of filters is initialised to 128 followed by 3 CNN layers of 60 filters each. The filter distribution for a similar experiment (60 filters on the first layer) is illustrated on figure 1. We observe that some filters with

Table 1: Filter pruning experiment: numbers of narrow band filters and related PER in function of the initialization.

Sinc-Layer num. filters	CNN-layers			narrow band filters	PER (%)
128	60	60	60	39	17.1
100	60	60	60	45	17.1
80	60	60	60	38	17.2
60	60	60	60	32	17.4
40	60	60	60	27	17.5
30	60	60	60	24	17.5

Table 2: SincNet experiment: compare the performance of the training with the filters fixed and the filters that are free to train.

num. filters	fixed filters		trained filters	
	loss	PER	loss	PER
40	2.35	18.3	2.31	17.6
30	2.37	18.0	2.33	17.5

a comparatively narrow bandwidth train towards a filterbank. The others train towards wider bandwidth filters. Concerning the frequency range, the wide-band filters could in principle be reconstructed with a linear combination of narrow band filters. In this paper, those two types of filters will be referred to as narrow and wide-band filters. The first part of this study focuses on the narrow band filters, since a large number of the wide-band filters seem to overfit the data.

3.3. Number of filters

Some filters in the first convolutional layer stay narrow-band while the others train towards wider bandwidths. Table 1 gives an overview of the number of narrow band filters as well as the phone error rate (PER) on TIMIT. To determine the number of narrow band filters an ad-hoc pruning operation has been applied after the filter training: the filters with wide bandwidths covering parts of the spectrum that are already covered by smaller bandwidth filters are discarded and only one filter is taken into account around the Nyquist frequency.

The number of filters needed by the model to build a filterbank covering the whole frequency range can be determined by the number of narrow band filters. From table 1 we can deduce that above 30 filters, the number of narrow-band filters that describe the frequency range is around 30 - 40 filters, this correlates with the results obtained by Zeghidour [24] using Gabor filters.

We also notice that when the first layer is initialized to 30 or 40 filters (corresponding to the number of narrow-band filters of other layers), some of those filters still train toward larger band filters. To analyse whether keeping the initialization to the initial scale performs as well as the combination of narrow and wide-band filters that the model learns, experiments have been made on 30 and 40 filters for fixed and non-fixed filters, the results are given in table 2. This raises the hypothesis that the wide-band filters are bringing some information not provided by the narrow band filters.

Table 3: Mean Euclidean distance between narrow bandfilter’s normalized central frequency distribution and different scalings for different amount of filters (Mel filterbank) and different initial scalings (30 filters).

Initialized to scale - filters	Compared to			
	Mel	Bark	ERB	Greenwood
Mel - 128	0.0023	0.0047	0.0070	0.0086
Mel - 60	0.0018	0.0044	0.0070	0.0088
Mel - 40	0.0022	0.0039	0.0065	0.0084
Mel - 30	0.0020	0.0043	0.0071	0.0091
Bark - 30	0.0025	0.0037	0.0062	0.0082
ERB - 30	0.0030	0.0029	0.0055	0.0076
Greenwood -30	0.0037	0.0068	0.0095	0.0116

3.4. Scale after training

Given that there are several frequency warpings that can be justified from a physiological point of view, it is informative (and simple) to investigate which one is preferred by an E2E system. It is clear by inspection that it is the narrow band filters that learn the warped filterbank. In [29] a convolutional variational autoencoder (CVAE) architecture that learns convolutional filters from raw waveforms using unsupervised learning was proposed. However, the analysis was only based on the central frequencies learned by those filters, not the narrow/wide-band distinction. In the present paper the central frequencies of only the narrow band filters are taken into account.

3.4.1. Experiment

The experiment consisted simply of analysing which filterbank the narrow-band filters trained above were learning. The experiment was repeated for several models with different initializations. The metric used to compute the distance between the initial and trained scale is the mean of the Euclidean distance:

$$d(x, s) = \frac{1}{N} \sqrt{\sum_{i=0}^N (x_i - s_i)^2} \quad (2)$$

The narrow band filters of a filterbank initialized to the mel scale remain mel-scale distributed. When initializing 30 – 40 filters as starting point to different scalings, other scalings also train towards mel scale. It follows that the mel scaling is an appropriate choice for filterbank initializations.

3.5. Corollary

The results so far have reinforced that E2E approaches do indeed learn what has been known for many years about cochlear models: that 30 to 40 filters are sufficient and that, regardless of physical measurements, the mel scale is the one that is perceptually important. However, from figure 1 it is clear that SincNet filters train towards a mixture of narrow *and* wide-band filters. Moreover from table 1, in all the experiments done for this section, the model always learns wide-band filters. It follows that these wide-band structures are important. Two questions arise:

1. Can the filters be initialized to wide-band, removing or reducing the need to train them?
2. Why do wide-band filters appear at all given that they are, to a first approximation, just linear combinations of narrow-band filters?

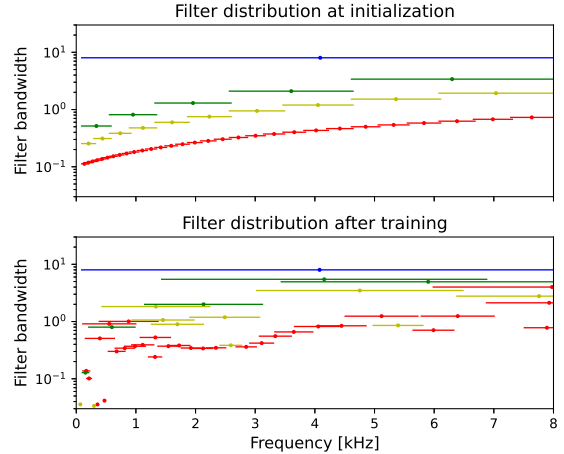


Figure 2: Filter repartition of superimposed filterbanks before (top plot) and after (bottom plot) training. In the initialization, the red filterbank is a narrow-band filterbank composed of 30 filters. The rest are filterbanks of 10, 5 and 1 filters capturing information that could in principle be reconstructed by a combination of the narrow band filters.

Table 4: Summary of experiments using narrow and/or wide-band filters

filters	filter type	fixed	loss	PER
10 - 5 - 1	wide	yes	2.410	18.6%
30	narrow	yes	2.374	18.0%
30 - 10 - 5 - 1	narrow & wide	yes	2.335	17.7%
30 - 10 - 5 - 1	narrow & wide	no	2.306	17.5%

These are addressed in the following section.

4. Wide-band filter analysis

4.1. Wide-band initialization

4.1.1. Hypothesis

Wide-band filters are important; it follows that the training can be optimized by initializing a narrow band filterbank as before and adding wide-band filters in addition of those filters. This hypothesis can be confirmed by an experiment comprising initializing several frozen superimposed filterbanks where the wide-band filters are combinations of several narrow band filters.

4.1.2. Experiment

Figure 2 shows the initialization and training of a model built with 4 ranges of filters illustrated on the upper plot. An estimation of the filter distribution after training of this new initialization is illustrated in the bottom plot. Four experiments using those filters are summarized in table 4:

Using only the narrow band filters gives a final PER of 18%, the combination of narrow and wide-band filters give for frozen filters a PER of 17.7% and for trained filters a PER of 17.5%. A combination of narrow and wide-band frozen filters already gives an improvement of 0.3% WER. The same effect is observed on the loss: the loss for a combination of narrow

and wide-band filters is lower than for only narrow band filters. The new initialization is consequently closer to what the model learns compared to the baseline initialization. Aside, it is interesting to notice from figure 2 that when trained most of the narrow band filters stay narrow band and most of the wide-band filters stay wide-band.

Thus the hypothesis is demonstrated. We conclude that it can indeed be beneficial to provide a mixture of narrow and wide-band filters in an ASR front-end.

4.2. Why wide-band filters?

Wide-band filters are in principle linear combinations of several narrow band filters; the network should be able to learn such a combination trivially, much as we assume it is learning the energy feature that was commonly used in HMM-based models. The most plausible explanation for the network’s failure to do so arises from the interaction of harmonic (voiced) and aperiodic (unvoiced) components. Harmonic components in the same filter add constructively in the magnitude domain. Aperiodic components, however, add as random variables; the variances add leading to a magnitude reduction by a factor of \sqrt{N} for N discrete components. The wider band filters hence tend to favour the voiced components.

In complex narrowband (e.g., Fourier transform based) filters, the squaring operation leads to both harmonic and aperiodic components adding *in the same ratio* in the magnitude or power domain, inhibiting emphasis of the harmonic component. SincNet comprises real-valued filters; however, the subsequent network architecture can inhibit the behaviour. Each convolutional layer is followed by four typical operations: max-pooling (downsampling), layernorm, ReLU (activation function) and dropout that have some influence on the signal. Of these, the maxpooling function and ReLU activation bring some distortions to the signal.

4.2.1. Hypothesis

It is possible to design a simple experiment to examine whether the above non-linearities inhibit simulation of wideband filters. The experiment encompasses three intuitive hypotheses:

1. **Using average-pooling instead of maxpooling** removes the noise artifacts that are created by maxpooling on the filtered signal, but since we continue to use a pooling function, aliasing still happens for the high frequencies. In [30] some experiments showed that using average pooling reduced the WER but without explaining the possible reasons.
2. **Moving the first downsampling factor towards a further layer** inhibits downsampling just after the filtering of the signal, this removes both the effect of aliasing and signal distortion (although it increases the data size at several convolution layers).
3. By inspection, **using a tanh or sigmoid function** removes some low frequency artifacts created by the ReLU function. However, it is well known that the cochlea contains a rectifier function [31], implying that ReLU is the right physiologically plausible solution. It is not clear a-priori which of these properties is more important.

4.2.2. Experiment

Table 5 summarises the performance of the experiments implied above after the first convolution layer: use average pooling, move the first downsampling factor to a later layer and check that ReLU is appropriate.

Replacing the max-pooling with average-pooling leads to

Table 5: *The effects of modifying the downsampling and pooling schemes. The numbers in the second column refer to the downsampling rate at each of the pooling operations in the convolutional layers (1 implies no downsampling).*

filters	downsampling	pooling 1st layer	act. function	PER
30	3-3-3-2	max	ReLU	17.5%
30	3-3-3-2	avg	ReLU	17.1%
30	1-3-3-6	-	ReLU	16.8%
30	1-3-3-6	-	sigmoid	18.1%

an improvement in performance. Displacing the downsampling by one layer, in principle allowing the network to combine filters, leads to a further improvement. This broadly demonstrates the first two points of our hypothesis. Changing the activation function to sigmoid deteriorates the PER. This tends to confirm that the physiologically-implied rectifier — yielding a simple spectral envelope — is also the right solution in the artificial solution. We note that, even with the best performing architecture, the system still learns some wide-band filters. This implies that our solution is not perfect.

5. Conclusion

E2E training of filterbanks for ASR leads to filters that resemble a standard filterbank. However, wider bandwidth filters are learned too, and are important for good ASR performance. The central frequencies of the narrow-band filters tend to a mel spacing, regardless of the initialisation. This confirms a well understood mechanism, suggesting that it exists in the biological system. There appears to be an optimal number of filters — around 30 to 40 — that also correlates with acknowledged literature.

We suggest that wide-band filters are learned to distinguish voiced (harmonic, coherent) components from either background noise or unvoiced (aperiodic, stochastic) components. In principle, a network should be able to learn such wide-band components by combining narrow-band ones. We argue that this capability is precluded by the (otherwise standard) ML architecture; in particular the phase will be lost by maxpooling. This argument is borne out by experiments showing that a structure retaining a more formal down-sampling mechanism can lead to better performance.

We are not aware of structures in the inner-ear or cochlea that can emulate physical wide-band filters. However the phase information is retained in our current best understanding of cochlear operation, retaining also the possibility that such filters are emulated in mid-brain. Proving this would be difficult, perhaps requiring some combination of selective stimuli with MRI or EEG sensing. It remains as a hypothesis for the neuroscience community. Our own future work will focus on more biologically plausible architectures. This experimental study indicates that any such model will need to retain phase in order for the subsequent network to take advantage of both narrow-band and wide-band features.

6. Acknowledgements

This work was supported as a part of NCCR Evolving Language, a National Centre of Competence in Research, funded by the Swiss National Science Foundation (grant number 180888).

7. References

- [1] D. Amodei, S. Ananthanarayanan, R. Anubhai, J. Bai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, Q. Cheng, G. Chen *et al.*, “Deep speech 2: End-to-end speech recognition in English and Mandarin,” in *International conference on machine learning*, PMLR, 2016, pp. 173–182.
- [2] S. Schneider, A. Baevski, R. Collobert, and M. Auli, “wav2vec: Unsupervised pre-training for speech recognition,” *arXiv preprint arXiv:1904.05862*, 2019.
- [3] A. Babu, C. Wang, A. Tjandra, K. Lakhota, Q. Xu, N. Goyal, K. Singh, P. von Platen, Y. Saraf, J. Pino *et al.*, “Xls-r: Self-supervised cross-lingual speech representation learning at scale,” *arXiv preprint arXiv:2111.09296*, 2021.
- [4] R. F. Lyon, *Human and Machine Hearing: Extracting Meaning from Sound*. Cambridge University Press, 2017.
- [5] B.-H. Juang and L. R. Rabiner, “Automatic speech recognition—a brief history of the technology development,” *Georgia Institute of Technology. Atlanta Rutgers University and the University of California. Santa Barbara*, vol. 1, p. 67, 2005.
- [6] H. Hermansky, “Perceptual linear predictive (plp) analysis of speech,” *the Journal of the Acoustical Society of America*, vol. 87, no. 4, pp. 1738–1752, 1990.
- [7] R. Collobert, C. Puhersch, and G. Synnaeve, “Wav2letter: an end-to-end convnet-based speech recognition system,” *arXiv preprint arXiv:1609.03193*, 2016.
- [8] M. Ravanelli and Y. Bengio, “Speaker recognition from raw waveform with sincnet,” in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 1021–1028.
- [9] G. Von Békésy, “Experiments in hearing,” *McGraw-Hill*, 1960.
- [10] C. D. Geisler, “Mathematical models of the mechanics of the inner ear,” in *Auditory system*. Springer, 1976, pp. 391–415.
- [11] J. Zwislocki, “Review of recent mathematical theories of cochlear dynamics,” *The Journal of the Acoustical Society of America*, vol. 25, no. 4, pp. 743–751, 1953.
- [12] E. De Boer, “On active and passive cochlear models—toward a generalized analysis,” *The Journal of the Acoustical Society of America*, vol. 73, no. 2, pp. 574–576, 1983.
- [13] A. Hudspeth, “Making an effort to listen: mechanical amplification in the ear,” *Neuron*, vol. 59, no. 4, pp. 530–545, 2008.
- [14] A. Hudspeth, F. Jülicher, and P. Martin, “A critique of the critical cochlea: Hopf—a bifurcation—is better than none,” *Journal of neurophysiology*, vol. 104 3, pp. 1219–29, 2010.
- [15] D. Sridhar, O. Stakhovskaya, and P. A. Leake, “A frequency-position function for the human cochlear spiral ganglion,” *Audiology and Neurotology*, vol. 11, no. Suppl. 1, pp. 16–20, 2006.
- [16] S. S. Stevens, J. Volkman, and E. B. Newman, “A scale for the measurement of the psychological magnitude pitch,” *The journal of the acoustical society of america*, vol. 8, no. 3, pp. 185–190, 1937.
- [17] E. Zwicker, “Subdivision of the audible frequency range into critical bands (frequenzgruppen),” *Journal of the Acoustical Society of America*, vol. 33, no. 2, p. 248, Feb. 1961.
- [18] J. O. Smith and J. S. Abel, “Bark and erb bilinear transforms,” *IEEE Transactions on speech and Audio Processing*, vol. 7, no. 6, pp. 697–708, 1999.
- [19] B. C. Moore and B. R. Glasberg, “Suggested formulae for calculating auditory-filter bandwidths and excitation patterns,” *The journal of the acoustical society of America*, vol. 74, no. 3, pp. 750–753, 1983.
- [20] B. J. Shannon and K. K. Paliwal, “A comparative study of filter bank spacing for speech recognition,” in *Microelectronic engineering research conference*, vol. 41. Citeseer, 2003, pp. 310–12.
- [21] L. Burget and H. Heřmanský, “Data driven design of filter bank for speech recognition,” in *International Conference on Text, Speech and Dialogue*. Springer, 2001, pp. 299–304.
- [22] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath *et al.*, “Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups,” *IEEE Signal processing magazine*, vol. 29, no. 6, pp. 82–97, 2012.
- [23] D. Palaz, R. Collobert, and M. M. Doss, “Estimating phoneme class conditional probabilities from raw speech signal using convolutional neural networks,” *arXiv preprint arXiv:1304.1018*, 2013.
- [24] N. Zeghidour, N. Usunier, I. Kokkinos, T. Schaiz, G. Synnaeve, and E. Dupoux, “Learning filterbanks from raw speech for phone recognition,” in *2018 IEEE international conference on acoustics, speech and signal Processing (ICASSP)*. IEEE, 2018, pp. 5509–5513.
- [25] M. Ravanelli and Y. Bengio, “Speech and speaker recognition from raw waveform with sincnet,” *arXiv preprint arXiv:1812.05920*, 2018.
- [26] H. Seki, K. Yamamoto, T. Akiba, and S. Nakagawa, “Discriminative learning of filterbank layer within deep neural network based speech recognition for speaker adaptation,” *IEICE TRANSACTIONS on Information and Systems*, vol. 102, no. 2, pp. 364–374, 2019.
- [27] J. S. Garofolo, “Timit acoustic phonetic continuous speech corpus,” *Linguistic Data Consortium, 1993*, 1993.
- [28] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [29] P. Agrawal and S. Ganapathy, “Unsupervised raw waveform representation learning for asr,” in *INTERSPEECH*, 2019, pp. 3451–3455.
- [30] H. Dubey, A. Sangwan, and J. H. Hansen, “Transfer learning using raw waveform sincnet for robust speaker diarization,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6296–6300.
- [31] E. De Boer and H. De Jongh, “On cochlear encoding: Potentialities and limitations of the reverse-correlation technique,” *The Journal of the Acoustical Society of America*, vol. 63, no. 1, pp. 115–135, 1978.