



Radically Old Way of Computing Spectra: Applications in End-to-End ASR

Samik Sadhu¹, Hynek Hermansky^{1,2}

¹Center for Language and Speech Processing, Johns Hopkins University, USA

²Human Language Technology Center of Excellence, Johns Hopkins University, USA

samiksadhu@jhu.edu, hynek@jhu.edu

Abstract

We propose a technique to compute spectrograms using Frequency Domain Linear Prediction (FDLP) that uses all-pole models to fit the squared Hilbert envelope of speech in different frequency sub-bands. The spectrogram of a complete speech utterance is computed by overlap-add of contiguous all-pole model responses. A long context window of 1.5 seconds allows us to capture the low frequency temporal modulations of speech in the spectrogram. For an end-to-end automatic speech recognition task, the FDLP spectrogram performs on par with the standard mel spectrogram features for clean read speech training and test data. For more realistic speech data with train-test domain mismatches or reverberations, FDLP spectrogram shows up to 25% and 22% relative WER improvements over mel spectrogram respectively.

Index Terms: Frequency Domain Linear Prediction, End-to-end Automatic Speech Recognition

1. Introduction

Since the advent of digital signal processing in speech, speech spectra for automatic speech recognition (ASR) are computed through short time analysis of speech from 10-20 msec speech segments. The spectral dynamics are provided through dynamic features or by concatenation of short-time spectral vectors over appropriate time spans. However, the original spectrograms [1] were derived from energies at outputs of band-pass filters covering the speech spectrum. We return to this original concept of spectrogram and use long term temporal analysis [2, 3, 4, 5, 6] to directly derive temporal modulations over long segments of speech.

Amongst the latter models, Frequency Domain Linear Prediction (FDLP) [7, 3, 8] is a technique to fit all-pole models to the squared Hilbert envelope of speech with varied degrees of approximation given by the model order leaving behind a frequency modulated component as residual [9, 10]. Firstly, the FDLP model shows similar “peak-hugging” characteristics like its more well-known dual, Time Domain Linear Prediction (TDLP) [11] and prioritizes high energy regions of the Hilbert envelope. Secondly, the all-pole approximation of the Hilbert envelope provides a straight-forward way to compute the rate of change of energy with time or *modulation spectrum* of speech. This can be done recursively from the autoregressive coefficients of the all-pole model [12] and allows for selective alleviation of some modulations from the model response when computing the FDLP spectrogram.

In section 2 we describe our speech processing technique to obtain the FDLP spectrogram. Subsequently, section 5 analyzes the results from end-to-end ASR models trained with

FDLP spectrogram and compares them with the traditional mel spectrogram features.

2. FDLP spectrogram

2.1. Frequency Domain Linear Prediction (FDLP)

Given samples of a signal x , the squared Hilbert envelope H of x is computed as the squared magnitude of the discrete time analytical signal of x [13]. It has been shown that linear prediction analysis of the Discrete Cosine Transform (DCT) of x yields an all-pole model which approximates the squared Hilbert envelope H with a degree of approximation given by the model order p [14].

In linear prediction analysis [11], Levinson-Durbin recursion can be used to obtain the model coefficients $\alpha_m, m = 1, 2, \dots, p$ for any specified model order p and the model gain G . We define the *FDLP response* F as the Fourier transform of the inverse of this model. Figure 1 shows how the FDLP response fits the energy of the signal x .




Figure 1: *FDLP response of a sample speech signal using an all-pole model of order 150*

2.2. Critical-band spectral trajectories using FDLP

Frequency bands can be formed by windowing of the frequency-domain DCT projection of x as described in the Section 2.6.

2.3. FDLP to modulation spectrum

The modulation spectrum captures the variations of the logarithmic energy of the signal x with time [15]. Given that the FDLP response approximates the squared Hilbert envelope H , which in turn represents the energy profile of x as a function of time, one reasonable definition of the modulation spectrum of x which we adopt in this work would be as given below.

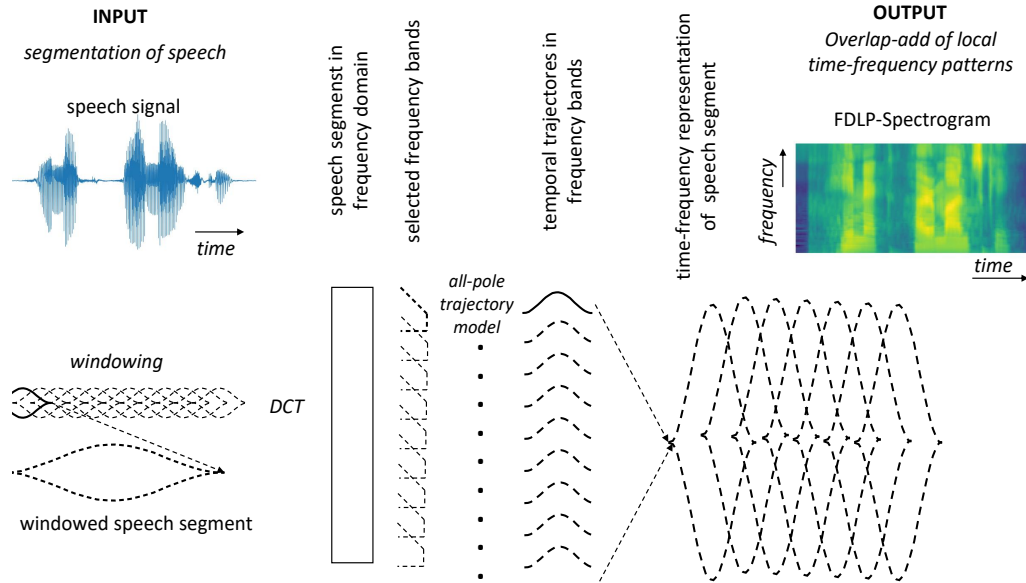


Figure 2: Speech signal is divided into uniform-length shorter speech segments. Each segment is weighted by von Hann window. The windowed segment is transformed to frequency domain through discrete cosine transform (DCT). Auditory-like weights are applied to appropriate parts of the DCT-transformed signal to provide inputs to all-pole modeling. The all-pole models approximate squared Hilbert envelopes of speech segments at different frequencies to yield the local time-frequency pattern of windowed speech segments. Overlap-add technique then yields the total time-frequency pattern of the input speech signal.

Definition 1 The modulation spectrum of the FDLP response is given by $\mathcal{M} = IDTFFT(\log F)$, where $IDTFFT$ is the inverse Discrete Time Fourier Transform.

Since linear prediction is guaranteed to give stable minimum phase systems, the modulation spectrum can also be computed by recursion directly from the all-pole model coefficients [16] as in eq. 1

$$\mathcal{M}[m] = \begin{cases} 0 & \text{for } m < 0 \\ \log G & \text{for } m = 0 \\ \alpha_m + \sum_{i=1}^{m-1} \frac{i}{m} \alpha_{m-i} \mathcal{M}[i] & \text{for } m > 0 \end{cases} \quad (1)$$

where $\alpha_m = 0$ for $m > p$.

For a T seconds long signal, \mathcal{M} is obtained as projections on cosines that are $\frac{1}{2T}$ Hz apart. Hence modulations upto F_m Hz are captured by the first $2F_m T$ coefficients in \mathcal{M} .

2.4. Modifying the FDLP response

Modulation frequencies relevant to speech recognition [17] are selected by applying weights $\gamma[m]$, $m = 1, 2, \dots, 2F_m T$ on different cosine projections. γ is similar to *cepstral liftering* of TDLP model since the FDLP-derived modulation spectrum is dual to the cepstrum of the TDLP all-pole model. The weighted modulations can be changed to a modified FDLP response as

$$\hat{F} = \exp(DTFT(\mathcal{M} \odot \gamma)) \quad (2)$$

where \odot is the point-wise multiplication operator and $DTFT$ is the Discrete Time Fourier Transform.

2.5. Windowing and overlap-add

The FDLP responses are computed over a fixed time duration of T seconds. However, speech utterances can be of variable duration. The T seconds segments of speech signal are weighted by cosine (von Hann) windows. This allows for applying quarter-window-length Overlap-Add (OLA) [18] of the FDLP response to concatenate the FDLP responses from the individual speech segment. This operation yields continuous temporal trajectory estimates of the whole speech utterance. The window also de-emphasizes less reliable end-points of the FDLP all-pole approximation.

2.6. FDLP spectrogram

The FDLP spectrogram is computed from FDLP responses in different frequency sub-bands for a given speech signal x . We use 80 cochlear filters [19] equally spaced in the bark scale to separate $DCT(x)$ into frequency sub-bands by point-wise multiplication. To capture low frequency temporal modulations in speech, we use long 1.5 second windows of speech for all-pole model estimation based on optimal clean development set ASR performance. Assuming a 100 Hz frame rate requirement for the ASR task, the FDLP spectrogram is computed as follows (see figure 2)

1. Window x using $T = 1.5$ seconds long von Hann windows with 25% overlap
2. For each windowed signal x_w , compute $D_w = DCT(x_w)$
3. Point-wise multiply D_w with 80 cochlear filter weights to obtain $D_w^{(1)}, D_w^{(2)}, \dots, D_w^{(80)}$.

4. Do linear predictive analysis of $D_w^{(1)}, D_w^{(2)}, \dots, D_w^{(80)}$.
5. Compute modulation spectrum $\mathcal{M}_w^{(1)}, \mathcal{M}_w^{(2)}, \dots, \mathcal{M}_w^{(80)}$ from each of the 80 set of linear prediction coefficients using the recursive formulation.
6. Apply weights γ on each modulation spectrum
7. Compute log FDLP responses from the weighted modulation spectrum down-sampled to the appropriate frame-rate of 100 Hz.
8. The spectrogram for the windowed speech x_w is obtained by forming a $80 \times 100T$ dimensional matrix of the FDLP responses.
9. The spectrogram of the complete signal x is computed by OLA of the spectrograms of the time-shifted windows.

The FDLP spectrogram looks similar to mel spectrogram even though the two spectrograms are computed by dual spectro-temporal processing techniques. However, the FDLP spectrogram **a)** focuses on capturing only energy concentrates in the Hilbert envelope and **b)** has the added flexibility of choosing different levels of robustness using the all-pole model order and **c)** manipulating the modulation spectrum. An implementation of FDLP spectrogram is available at https://github.com/sadhusamik/speech_recognition_tools.

3. Mel spectrogram

The baseline mel spectrogram (also referred to as Log Filter Bank Energy) features are obtained by short-time analysis of the signal x with 20 ms Hamming windows and a frame-rate of 100 Hz. We compute the magnitude spectrum for each windowed signal. The log spectral energy in 80 mel-scaled triangular filters applied to the magnitude spectrum generates a 80 dimensional vector every 10 ms. These vectors are concatenated over one speech utterance to obtain the mel spectrogram.

4. Experimental Setup

4.1. FDLP spectrogram configuration

4.1.1. Window length

We use $T = 1.5$ seconds long von Hann windows to compute the FDLP response.

4.1.2. Model order

A model order of p allows the all-pole model to fit a maximum of $\lfloor \frac{p}{2} \rfloor$ energy peaks of the Hilbert envelope. In section 5 we show how the ASR performance varies with the model order.

4.1.3. Liftering

In this work, we only use binary lifters of the form $\gamma[m; a, b]$, where

$$\gamma[m; a, b] = \begin{cases} 1 & \text{if } a \leq m \leq b \\ 0 & \text{else} \end{cases} \quad (3)$$

Hence, for a window length $T = 1.5$ seconds, a lifter $\gamma[m; 0, 150]$ completely eliminates any cosine projections above $\frac{150}{2 \times 1.5} = 50$ Hz. Whereas, to eliminate the DC projection, we can apply a lifter $\gamma[m; 1, 150]$.

4.2. End-to-end ASR model

We use the standard transformer based `espnet1` end-to-end model recipe in the ESPnet [20] speech recognition toolkit which uses a joint attention-CTC [21] multi-task learning neural network setup. Experiments are done with 12 layers and 6 layers of encoder and decoder respectively with 2048 hidden nodes. A RNN language model is used along-side the acoustic model as in the standard `espnet1` recipe. The specificities of the training configuration can be found in https://github.com/sadhusamik/speech_recognition_tools/blob/master/e2e/reverb/conf/train.yaml

4.3. Data sets

To analyze the performance of FDLP spectrogram for clean read speech as well as more practical reverberated speech and domain mismatched conditions, we train different ASR models on the following data sets

4.3.1. WSJ: clean read speech

An ASR model is trained with the whole of `si_284` data from Wall Street Journal (WSJ) consisting of 73 hours of labelled clean read speech. The model is tested with the clean test set `test_eval92`, as well as two artificially corrupted test sets generated by using 20dB of additive street noise and babble noise on `test_eval92` respectively. These two additional test sets are named `street20` and `babble20` respectively.

4.3.2. REVERB: noisy reverberated speech

The performance of FDLP spectrogram on reverberated speech data is evaluated by training an ASR model with *simulated* 8 channel reverberated data from the REVERB challenge [22]. As in the standard ESPnet recipe, this training data is augmented with clean speech from the `si_284` training set of WSJ. We test the model on three test sets consisting of *real* reverberated speech data, namely **a)** `real_1ch`: 1 channel speech data with no pre-processing, **b)** `real_1ch_wpe`: 1 channel speech data with WPE de-reverberation [23], **c)** `real_8ch`: 8 channel speech data with WPE de-reverberation and beamforming (using `beamformIt` [24]).

For both the data sets, we do not use additional data augmentation techniques (like speed perturbation [25] and SpecAugment [26]) which have been used in recent literature to boost end-to-end performance with ESPnet [27].

5. Results

5.1. Results on WSJ

Table 1 shows how the ASR performance on the clean test set `test_eval92` varies with changing model order and lifter configuration. It can be seen that model orders higher than 150 does not add any significant gain to the ASR performance. Previous studies have observed that modulation frequencies in 1-16 Hz range are the most important for ASR as well has human speech cognition [17]. In our experiments, we observed noticeable improvements by including cosine projections till 33 Hz (see table 1). However, addition of further modulations adversely affects the ASR performance. Leaving out the DC component of the modulation spectrum slightly degrades performance on this acoustically well controlled WSJ data.

Table 2 shows a comparison of published state-of-the-art ES-
Table 1: Performance on *test_eval192* with (a) modulation range 0-33 Hz and various model orders, (b) model order 150 and various lifter configurations

(a)		(b)	
model order (p)	WER %	lifter configuration (γ)	WER %
80	5.7	a=0, b=75	5.5
100	5.3	a=0, b=100	4.8
150	4.8	a=0, b=150	5.2
200	4.8	a=0, b=300	5.1
		a=0, b=450	5.0
		a=1, b=100	5.3

Pnet performances on WSJ using the same model architecture, our implementation of mel spectrogram and FDLP spectrogram with $p = 150$, and modulations in the range 0-33 Hz. It can be seen that FDLP spectrogram performs at-par with the state-of-the-art ESPnet models using mel spectrogram+pitch features. In addition, FDLP spectrogram shows significantly better performance on the noisy mismatched test sets *street20* and *babble20* with up to 25% relative WER improvements.

The constant energy in each frequency sub-bands, represented by the DC component of the modulation spectrum, can easily get corrupted by slowly changing environmental factors or any variations in the microphone characteristics [28]. We explore the effects of removing lower cosine projections from our modulation spectrum under reverberated speech conditions in section 5.2. However, for clean read speech, removing the DC projection (see table 1(b)) can marginally affect ASR performance.

Table 2: Comparison of mel spectrogram and FDLP spectrogram performance on WSJ

Features	WER %		
	<i>test_eval192</i>	<i>street20</i>	<i>babble20</i>
Guo et al. [27] *	4.9	-	-
our mel spectrogram	5.1	24.7	75.2
FDLP spectrogram	4.8	20.4	56.1

* This result uses mel spectrogram + pitch as features as well as SpecAugment data augmentation technique.

5.2. Results on REVERB

The effect of reverberation is captured in low modulation frequencies. Table 3(a) shows the performance of the ASR model on the *real_8ch* test set trained with FDLP spectrogram using a model order of 150 and different ranges of low cosine projections removed.

On the other hand, table 3(b) shows that including higher cosine projections up to 450 Hz achieves a better performance. The reverberated signal is generated by convolving clean speech with room impulse responses that smooth out sudden transitions that characterize plosive sounds. Addition of these higher

Table 3: Performance on *real_8ch* with model order 150 as a result of (a) removing low modulations (b) including higher modulations from FDLP spectrogram

(a)		(b)	
lifter configuration (γ)	WER %	lifter configuration (γ)	WER %
a=0, b=100	8.5	a=1, b=75	8.4
a=1, b=100	7.9	a=1, b=100	7.9
a=2, b=100	8.0	a=1, b=150	7.7
		a=1, b=300	7.8
		a=1, b=450	7.2
		a=1, b=600	7.7

modulations better preserves remaining abrupt transitions in the FDLP spectrogram. In fact, in our experiments, plosive characters like B and P show up to 5% reduction in recognition error when modulations between 100 to 150 Hz are included in the FDLP spectrogram.

Table 4: Comparison of mel spectrogram and FDLP spectrogram performance together with results from recent literature on REVERB data.

Features	WER %		
	<i>real_8ch</i>	<i>real_1ch</i>	<i>real_1ch_wpe</i>
Guo et al. [27] *	14.3	-	-
Zhang et al. [29] †	10.0	-	-
our mel spectrogram	9.2	23.2	20.7
FDLP spectrogram	7.2	19.4	18.0

* This result uses mel spectrogram + pitch as features as well as speed perturbation data augmentation technique.

† Unifies de-reverberation and ASR under one architecture.

Table 4 shows that FDLP spectrogram has a 22% relative WER improvement over our mel spectrogram features. Additionally, ASR performance using FDLP spectrogram *without* WPE front-end de-reverberation is better than mel spectrogram *with* WPE de-reverberation. Thus, FDLP spectrogram is more effective at dealing with the effects of reverberation as compared to WPE. Using WPE front-end processing as well as FDLP spectrogram features yield a 13% relative WER reduction over mel spectrogram with WPE de-reverberation.

6. Conclusions

In this work we described a way to compute spectrograms using Frequency Domain Linear Prediction with several robustness benefits. The proposed spectrogram shows significant improvements over mel spectrogram for domain mismatched train-test scenarios as well as noisy, reverberated speech data and is better at handling the effects of reverberation compared to WPE alone.

7. Acknowledgements

This work was funded by the faculty gift from Google Research and by support from the JHU Human Language Technology Center of Excellence.

8. References

- [1] W. Koenig, H. Dunn, and L. Lacy, "The sound spectrograph," *The Journal of the Acoustical Society of America*, vol. 18, no. 1, pp. 19–49, 1946.
- [2] H. Hermansky and S. Sharma, "Temporal patterns (TRAPS) in asr of noisy speech," in *1999 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings. ICASSP99 (Cat. No. 99CH36258)*, vol. 1. IEEE, 1999, pp. 289–292.
- [3] M. Athineos, H. Hermansky, and D. P. Ellis, "LP-TRAP: Linear predictive temporal patterns," 2004.
- [4] S. Thomas, S. Ganapathy, and H. Hermansky, "Recognition of reverberant speech using frequency domain linear prediction," *IEEE Signal Processing Letters*, vol. 15, pp. 681–684, 2008.
- [5] S. Sadhu, R. Li, and H. Hermansky, "M-vectors: sub-band based energy modulation features for multi-stream automatic speech recognition," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6545–6549.
- [6] S. Sadhu and H. Hermansky, "Modulation vectors as robust feature representation for ASR in domain mismatched conditions," in *INTERSPPEECH*, 2019, pp. 3441–3445.
- [7] J. Herre and J. D. Johnston, "Enhancing the performance of perceptual audio coders by using temporal noise shaping (TNS)," in *Audio Engineering Society Convention 101*. Audio Engineering Society, 1996.
- [8] R. R. Shenoy and C. S. Seelamantula, "Frequency domain linear prediction based on temporal analysis," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 2629–2633.
- [9] S. Ganapathy, S. Thomas, P. Motlicek, and H. Hermansky, "Applications of signal analysis using autoregressive models for amplitude modulation," in *2009 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*. IEEE, 2009, pp. 341–344.
- [10] R. Kumaresan and A. Rao, "Model-based approach to envelope and positive instantaneous frequency estimation of signals with speech applications," *The Journal of the Acoustical Society of America*, vol. 105, no. 3, pp. 1912–1924, 1999.
- [11] J. Makhoul, "Linear prediction: A tutorial review," *Proceedings of the IEEE*, vol. 63, no. 4, pp. 561–580, 1975.
- [12] A. V. Oppenheim and R. W. Schaffer, "From frequency to quefrency: A history of the cepstrum," *IEEE signal processing Magazine*, vol. 21, no. 5, pp. 95–106, 2004.
- [13] L. Marple, "Computing the discrete-time analytic signal via FFT," *IEEE Transactions on signal processing*, vol. 47, no. 9, pp. 2600–2603, 1999.
- [14] S. Ganapathy, "Signal analysis using autoregressive models of amplitude modulation," Ph.D. dissertation, Johns Hopkins University, 2012.
- [15] N. Ding, A. D. Patel, L. Chen, H. Butler, C. Luo, and D. Poeppel, "Temporal modulations in speech and music," *Neuroscience & Biobehavioral Reviews*, vol. 81, pp. 181–187, 2017.
- [16] A. Oppenheim and R. Schaffer, "Homomorphic analysis of speech," *IEEE Transactions on Audio and Electroacoustics*, vol. 16, no. 2, pp. 221–226, 1968.
- [17] H. Hermansky, "The modulation spectrum in the automatic recognition of speech," in *1997 IEEE Workshop on Automatic Speech Recognition and Understanding Proceedings*. IEEE, 1997, pp. 140–147.
- [18] J. Allen, "Short term spectral analysis, synthesis, and modification by discrete Fourier transform," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 25, no. 3, pp. 235–238, 1977.
- [19] 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.
- [20] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N.-E. Y. Soplin, J. Heymann, M. Wiesner, N. Chen *et al.*, "ES-Pnet: End-to-end speech processing toolkit," *Proc. Interspeech 2018*, pp. 2207–2211, 2018.
- [21] S. Kim, T. Hori, and S. Watanabe, "Joint CTC-attention based end-to-end speech recognition using multi-task learning," in *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 4835–4839.
- [22] K. Kinoshita, M. Delcroix, T. Yoshioka, T. Nakatani, E. Habets, R. Haeb-Umbach, V. Leutnant, A. Sehr, W. Kellermann, R. Maas *et al.*, "The REVERB challenge: A common evaluation framework for dereverberation and recognition of reverberant speech," in *2013 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*. IEEE, 2013, pp. 1–4.
- [23] T. Nakatani, T. Yoshioka, K. Kinoshita, M. Miyoshi, and B.-H. Juang, "Speech dereverberation based on variance-normalized delayed linear prediction," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 7, pp. 1717–1731, 2010.
- [24] X. Anguera, C. Wooters, and J. Hernando, "Acoustic beamforming for speaker diarization of meetings," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 7, pp. 2011–2022, 2007.
- [25] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, "Audio augmentation for speech recognition," in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [26] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," *Proc. Interspeech 2019*, pp. 2613–2617, 2019.
- [27] P. Guo, F. Boyer, X. Chang, T. Hayashi, Y. Higuchi, H. Inaguma, N. Kamo, C. Li, D. Garcia-Romero, J. Shi *et al.*, "Recent developments on ESPnet toolkit boosted by conformer," *arXiv preprint arXiv:2010.13956*, 2020.
- [28] H. Hermansky and N. Morgan, "RASTA processing of speech," *IEEE transactions on speech and audio processing*, vol. 2, no. 4, pp. 578–589, 1994.
- [29] W. Zhang, A. S. Subramanian, X. Chang, S. Watanabe, and Y. Qian, "End-to-end far-field speech recognition with unified dereverberation and beamforming," *Proc. Interspeech 2020*, pp. 324–328, 2020.