

# SPATIAL ATTENTION FOR FAR-FIELD SPEECH RECOGNITION WITH DEEP BEAMFORMING NEURAL NETWORKS

Weipeng He<sup>1,2</sup>, Lu Lu<sup>3</sup>, Biqiao Zhang<sup>3</sup>, Jay Mahadeokar<sup>3</sup>, Kaustubh Kalgaonkar<sup>3</sup>, Christian Fuegen<sup>3</sup>

<sup>1</sup>Idiap Research Institute, Switzerland

<sup>2</sup>École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

<sup>3</sup>Facebook, 1 Hacker Way, Menlo Park, CA 94025, USA

weipeng.he@idiap.ch lulu0314@gmail.com {didizbq, jaym, kaustubhk, fuegen}@fb.com

## ABSTRACT

In this paper, we introduce spatial attention for refining the information in multi-direction neural beamformer for far-field automatic speech recognition. Previous approaches of neural beamformers with multiple look directions, such as the factored complex linear projection, have shown promising results. However, the features extracted by such methods contain redundant information, as only the direction of the target speech is relevant. We propose using a spatial attention subnet to weigh the features from different directions, so that the subsequent acoustic model could focus on the most relevant features for the speech recognition. Our experimental results show that spatial attention achieves up to 9% relative word error rate improvement over methods without the attention.

**Index Terms**— Deep beamforming networks, multi-channel far-field speech recognition, array signal processing, attention

## 1. INTRODUCTION

Smart speakers have become increasingly popular in people’s daily life, and far-field Automatic Speech Recognition (ASR) is one of the important technologies behind them. Far-field ASR is more challenging than near-field ASR as the far-field signal is more corrupted by reverberation, background noise, as well as interfering voices. In this paper, we address the far-field ASR problem by introducing deep beamforming neural networks with spatial attention, which show improvements over existing methods.

Previous studies have applied traditional signal processing techniques to the corrupted signal before the ASR, such as noise reduction [1], dereverberation [2], beamforming [3,4], and post-filtering [5,6]. Among them, beamforming exploits the spatial information in the multi-channel signal, so that the sounds from different directions can be enhanced or suppressed. Technically, the beamformers apply filtering on the input and summation across channels. The filter coefficients (i.e., beamformer weights) are chosen by solving objective functions that make the enhanced signal better quality under certain assumptions. The performance of the traditional beamforming techniques is often limited in complex real-world environments. This is because that these methods often rely heavily on the assumptions of some ideal conditions about the environments, such as stationary signal, high signal-to-noise ratio (SNR), and precisely estimated direction of arrival or steering vectors. Furthermore, their objective functions are indirect with regard to ASR. Therefore, the enhanced signal does not necessarily improve the ASR results.

More recently, deep neural networks for beamforming have been shown to outperform traditional methods, as they do not require strong assumption about the environments. These deep beamforming network approaches may be divided into three main categories:

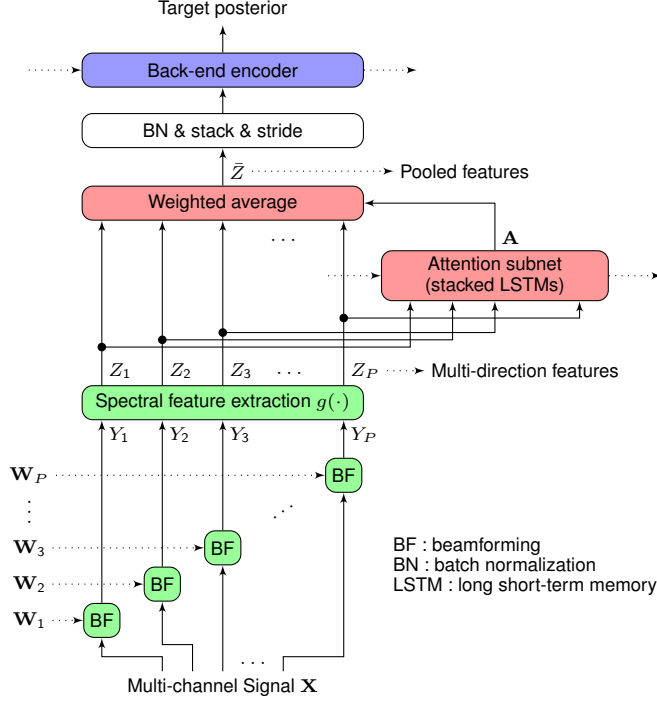
1. *Mask estimation* [7–9]: the neural network predicts the speech mask (and noise mask), which is then used for estimating the covariance matrix for Minimum Variance Distortionless Response (MVDR) or Generalized Eigenvector (GEV) beamforming.
2. *Adaptive filter coefficients estimation on single look direction* [9–13]: the neural network predicts the beamforming weights as an intermediate output. The weights enhance the signal in one look direction.
3. *Fixed filter coefficients estimation for beamforming on multiple directions* [14–16]: the beamforming weights are part of the neural network parameters. They are optimized to extract features on multiple directions and kept fixed after training.

Among these approaches, the factored complex linear projection (fCLP) from the third category has shown promising results with real-world smart speakers [16]. This approach is advantageous as it is computational efficient, does not require Direction of Arrival (DOA) estimation, and jointly optimizes enhancement and recognition. In fCLP, the features on all look directions are used as the input of the back-end network for Acoustic Modeling (AM). However, such features contain redundant information, as ideally only the features from the direction of the target speech are useful for the recognition.

In this paper, we propose using spatial attention to refine the multi-direction features in the fCLP approach. The spatial attention, computed from multi-directional features, indicates how informative each direction is for recognizing the target speech. We use attention score to weigh the features from the multiple directions using average pooling. This allows the subsequent acoustic model to focus on the features most relevant to ASR and reduces the number of parameters of the network, compared to the original multi-directional setting. The experiment results show that there is a significant improvement by adding spatial attention to the original fCLP approach.

Previous studies have explored similar ideas by using attention to select the most informative input channels for multi-channel far-field ASR [17, 18]. Our approach differs from them as we apply attention weighting to the spatial filters, not to the input channels.

\*This work was done when the first author was an intern with Facebook.



**Fig. 1.** Deep beamforming neural network with spatial attention, containing the neural beamformer (green), spatial attention subnet (red), and the back end (blue). Figure shows the processing of a single frame. The time dependency is indicated by the dotted arrows.

## 2. APPROACH

We propose an end-to-end neural network (Figure 1) for acoustic modeling from raw multi-channel signals with three components:

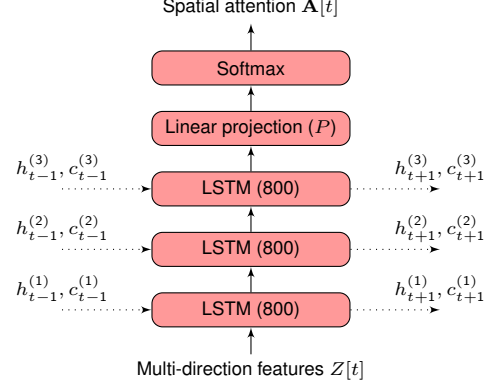
- *Neural beamformer*, which extracts speech features on multiple look directions.
- *Attention-based pooling module*, which picks the most informative features by spatial attention.
- *Back end*, which predicts the sequence of target posteriors (e.g., context-dependent phones or graphemes) from the sequence of pooled features.

The three components are trained jointly, so that the speech enhancement and feature extraction front-end are directly optimized to reduce the target classification error.

### 2.1. Front end: Neural Beamformer

The neural beamformer front end is based on the factored Complex Linear Projection (fCLP) [15, 16]. Specifically, the front-end takes multi-channel short-time Fourier transform (STFT)  $\mathbf{X}[t, f] = (X_1[t, f], \dots, X_M[t, f])^T \in \mathbb{C}^M$  as the network input.  $t$  and  $f$  are the indices of the time and frequency bin, respectively,  $M$  is the number of microphones, and  $^T$  is the matrix transpose. The input signal is beamformed with  $P$  spatial filters (look directions) using the weights  $\mathbf{W}_p[f] \in \mathbb{C}^M$ :

$$Y_p[t, f] = \mathbf{W}_p^H[f] \mathbf{X}[t, f], \quad p = 1, 2, \dots, P, \quad (1)$$



**Fig. 2.** The attention subnet.  $h$  and  $c$  are the hidden and cell states of the LSTMs. For online attention, the output  $\mathbf{A}[t]$  is smoothed by averaging over time (omitted in the figure). For offline attention, the output is taken at the last frame.

where  $p$  is the direction index and  $^H$  is the conjugate transpose. The beamforming weights  $\mathbf{W}_p[f]$  are parameters of the neural network.

After the spatial filtering (beamforming), we extract  $L$  spectral features from each enhanced signal  $Y_p$ :

$$Z_{p,l}[t] = g_l(Y_p[t]), \quad l = 1, 2, \dots, L, \quad (2)$$

where  $Y_p[t] = (Y_p[t, 1], \dots, Y_p[t, F])^T$ . In our experiments, we use the Complex Linear Projection (CLP) [19] as the spectral features extractor. That is:

$$Z_{p,l}[t] = \log \left| \sum_{f=1}^F Y_p[t, f] G_l[f] \right|, \quad (3)$$

where  $G_l = (G_l[1], \dots, G_l[F])^T \in \mathbb{C}^F$  are the parameters of the network. We have tested another feature extraction method with multi-layer perceptron in the preliminary experiments but did not find significant difference. The output of the beamformer front-end on each frame is a 2-D tensor with the dimensions of the directions and the spectral features, that is  $Z[t] = \{Z_{p,l}\} \in \mathbb{R}^{P \times L}$ .

Note that while the functions are written in complex form, the network computes the real and imaginary parts as separate real-valued functions, which allows us to calculate the derivative of the loss for back-propagation.

### 2.2. Spatial Attention

We use an attention-based pooling module to select the most informative direction. Given that the target speech is from a specific direction, only one or a few out of the  $P$  directions are relevant for recognizing the target speech, and the features from all other directions are redundant. Therefore, we add an subnet to decide how much attention should the ASR network pay to each direction. The output of the subnet, i.e., the spatial attention  $\mathbf{A}[t] = (A_1[t], \dots, A_P[t])^T \in [0, 1]^P$ , is used as the weight for the weighted average pooling across directions on the multi-direction features:

$$\bar{Z}_l[t] = \sum_{p=1}^P A_p[t] Z_{p,l}[t], \quad l = 1, 2, \dots, L. \quad (4)$$

The network predicts the spatial attention with stacked long short-term memory (LSTM) recurrent neural networks [20] (Figure 2). The hidden states are projected to  $P$  directions and normalized with a softmax layer.

Depending on when we take the output attention scores, we categorize the attention into *online*, *offline* and *offline with controlled latency*. The *online* version outputs attention scores for each direction at each frame. Considering that the direction of the sound source may not rapidly change, we apply averaging with a sliding window through frames to smooth the attention prediction. Whereas, in the *offline* version, the attention scores at the last frame of the sequence is applied to all frames. If we assume the sound source is not moving within one utterance, the whole utterance provides more information for the network to select the correct direction than just using local frames. The drawback of the offline version is the increased latency, which would not allow us to use this method for a streaming ASR. Therefore, we also proposed a *latency-controlled* version, which uses a short segment for attention prediction and apply the scores at a certain frame to the whole sequence, thus constraining the latency. The motivation of this approach is that when a trigger word is used to wake a device (e.g. “OK google”, “Alexa”, “Hey Portal”), the system could rely on the wake word segment to predict the attention without adding latency to the ASR.

### 2.3. Back end: Acoustic Model

After the attention-based pooling layer, the pooled features are stacked in 8 frames and subsampled by a factor of 3. The stacked features are then used as the input of the acoustic model, which predicts a sequence target posteriors of context-dependent phones or graphemes from the input feature sequence. Our approach suits arbitrary acoustic model neural networks. In this paper, we experiment with stacked uni-directional LSTMs as the back ends.

## 3. EXPERIMENTS

We compare the attention-based deep beamforming neural network for far-field ASR with a traditional signal processing-based front end as well as fCLP with other pooling methods.

### 3.1. Data

We use an in-house anonymized dataset collected through crowd-sourcing in the experiments. The crowd-sourced workers were asked to record commands for a smart assistant on mobile devices. Example commands include calling a contact, playing music, setting up timer/alarm, or getting time/weather information. The average duration of the utterances is around three seconds. More details about this dataset can be found in [21].

The training data are generated from these original near-field recordings in this dataset as follows: We randomly sample two million utterances (around 2000 hours) and simulate speech in reverberant and noisy environments. Specifically, for a given utterance, we first randomly simulate a room with a response time (RT60) between 200ms and 900ms, with different locations of the device microphone array, the speech source, and the noise source in the room. The average distance between the device microphone array and the speech source is three meters. The microphone geometry is set to match the device used to record the evaluation data: 4 microphones roughly forming a rectangle of  $6 \times 7$  cm. We then randomly sample a background noise segment from an in-house collected dataset.

Next, the clean near-field speech and the noise segment are both convolved with the simulated room impulse responses (RIR), and added together. We control the signal-to-noise ratio (SNR) of the resulting far-field noisy speech segment to be between 0 and 25 dB.

The evaluation data is collected by playing clean speech and noise from loudspeakers at various angles and distances, and recording with Portal+ devices. The clean speech and the noise are different segments sampled from the same datasets used for training. The evaluation set includes 16500 utterances (around 16 hours).

### 3.2. Implementation Details

The baseline methods includes the log-mel feature approaches (with or without a traditional Digital Signal Processing (DSP) front end) and the fCLP without spatial attention. The DSP front end baseline consists of the typical speech enhancement algorithms, including dereverberation, beamforming, and post-filtering. For the baseline method with no enhancement, the log-mel features are extracted from the first channel of the raw audio data (we can assume that which channel is selected does not affect the result). We compute 80 log-mel features with a window size of 25 ms and shift of 10 ms on the enhanced signal and use them as the input to the acoustic model.

For the fCLP-based models, the inputs are spectrograms with the same windowing as the log-mel features. The input is optionally dereverberated using the Weighted Prediction Error (WPE) algorithm [2]. We use the block-online implementations from the NARA-WPE library [22]. For the beamformer front end, we use  $P = 10$  look directions and  $L = 120$  spectral features. In addition to the attention-based pooling, we include the fCLP with no pooling (concatenate features from all directions [16]), max-pooling, and average-pooling (all directions with equal attention) for comparison.

In terms of the back end, we use a 5-layer uni-directional LSTM with 1200 units per layer for all approaches. We delay the prediction by 10 frames, so that there are more information available for the prediction. The output of the acoustic model is the posterior probability of 8576 context dependent graphemes [21]. We use the cross entropy as the loss function. The front and back ends are jointly trained with the adam optimizer [23] for 20 epochs. The learning rate is set to 0.001 and reduced by half whenever the validation loss does not decrease at the end of an epoch. During evaluation, we use the Weighted Finite-State Transducer (WFST) decoder [24] with a 4-gram language model.

## 4. RESULTS

We evaluate the performance of the proposed method and visualize its the directivity and attention patterns.

### 4.1. Performance Comparison

We compare different front-end approaches combined with the same back end (Table 1). The results show that the fCLP without pooling does not perform better than the log-mel approach without speech enhancement. This may be because of the model using all the features without pooling contains more parameters, and thus could be more prone to overfit. Adding the spatial attention let the fCLP outperform the traditional methods. Compared to no pooling, the online attention achieves 17% relative improvement if no dereverberation is applied. Applying dereverberation further improves the results. However, it also reduces the gap between the online attention and no pooling. The dereverberation reduces the potential mismatch be-

**Table 1.** Performance of the far-field ASR systems with the spatial attention compared with the baseline approaches. The dereverberation is optionally applied to both the training and evaluation data for the fCLP-based approaches.

Front-end Approach	WER (%)	
No enhancement, log-mel	14.3	
DSP, log-mel	13.8	
fCLP, with pooling	Dereverberation	
	None	Online WPE
None (baseline)	15.6	12.7
Max	13.8	12.8
Average	14.3	13.0
Attention (online)	13.0	12.4
Attention (latency 0.5s)	12.9	11.7
Attention (latency 1s)	12.6	11.5
Attention (offline)	12.6	11.3

tween the training and evaluation signals, hence mitigates the over-fitting issue of the no pooling approach.


The best performance is achieved by the offline attention, which uses whole utterances to predict the attention scores. However, the offline computation compromises the real-time responsiveness of the ASR system. As a trade-off, the attention with controlled latency uses half second or one second to predict the attention scores, performs almost as good as the offline attention. Compared to no pooling, the spatial attention with one second latency achieves 9% relative improvement. The results support our idea of potentially using trigger word for attention prediction.

## 4.2. Directivity Patterns


We plot the directivity patterns of the neural beamformers (Figure 3). The directivity patterns show how the signal from a direction is amplified (or attenuated). In contrary to what we expected, the directivity patterns show that the beamformers do not amplify on a single look direction (one red blob in the figure). Instead, each beamformer amplifies in many directions. Moreover, the gain between the amplified and attenuated directions is small (up to around 3 dB for white noise). This is likely due to that the number of spatial filters being much fewer than the possible signal directions. Therefore, the neural beamformer needs to cover the unit sphere with limited filters. Furthermore, because in our training data there are only one noise source in each utterance, a spatial filter is effective as long as it amplifies at the speech’s direction and attenuate in that noise direction, while all other directions do not matter.

## 4.3. Attention Patterns

We plot the predicted attention scores of the online spatial attention on some of the evaluation utterances (Figure 4). The patterns show that the attention of the first two seconds of an utterance is evenly distributed among all spatial filters (look directions). It is because the network requires certain amount of information to pick the correct filters. Unlike what we expected, even after the first two seconds, the attention is focused on a few filters instead of one. Empirically, the averaging of features from multiple directions helps the recognition (which also can be seen by comparing the performance between average-pooling vs no pooling). We have also found that there are



**Fig. 3.** Directivity patterns of the neural beamformers trained with online attention. Each figure shows the directivity pattern of one spatial filter. The directions on the unit sphere is projected into a grid of azimuth and elevation angles. The color shows the gain in dB with respect to the most amplified direction (assuming the signal is white noise). Similar directivity patterns have been found on the offline attention (which is therefore omitted).



**Fig. 4.** Visualization of the online spatial attention. It shows the attention score of each look direction (spatial filter) at each frame.

some directions constantly getting more attention than the others. Additionally, the number of directions saturates. This is consistent with our experiments that more look directions ( $P = 20$ ) does not improve the WER.

## 5. CONCLUSION

We have proposed spatial attention for multi-direction deep beamforming neural networks. The spatial attention is used to select the directions to attend to, so that it reduces the dimension of the neural beamformer features while keeping the most informative features of the target speech. Our experiments with far-field ASR on a smart speaker device show that the deep beamforming neural network with spatial attention performs up to 9% better than that without the attention.

## 6. REFERENCES

- [1] S. Boll, "Suppression of acoustic noise in speech using spectral subtraction," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 27, no. 2, pp. 113–120, Apr. 1979.
- [2] T. Yoshioka and T. Nakatani, "Generalization of Multi-Channel Linear Prediction Methods for Blind MIMO Impulse Response Shortening," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 10, pp. 2707–2720, Dec. 2012.
- [3] H. Cox, R. Zeskind, and M. Owen, "Robust adaptive beamforming," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 35, no. 10, pp. 1365–1376, Oct. 1987.
- [4] Joerg Bitzer and K. Uwe Simmer, "Superdirective Microphone Arrays," in *Microphone Arrays*, Arild Lacroix, Anastasios Venetsanopoulos, Michael Brandstein, and Darren Ward, Eds., pp. 19–38. Springer Berlin Heidelberg, Berlin, Heidelberg, 2001.
- [5] K. Uwe Simmer, Joerg Bitzer, and Claude Marro, "Post-Filtering Techniques," in *Microphone Arrays*, Arild Lacroix, Anastasios Venetsanopoulos, Michael Brandstein, and Darren Ward, Eds., pp. 39–60. Springer Berlin Heidelberg, Berlin, Heidelberg, 2001.
- [6] I. A. McCowan and H. Bourslard, "Microphone array post-filter based on noise field coherence," *IEEE Transactions on Speech and Audio Processing*, vol. 11, no. 6, pp. 709–716, Nov. 2003.
- [7] Hakan Erdogan, Tomoki Hayashi, John R Hershey, Takaaki Hori, Chiori Hori, Wei-Ning Hsu, Suyoun Kim, Jonathan Le Roux, Zhong Meng, and Shinji Watanabe, "Multi-Channel Speech Recognition: LSTMs All the Way Through," in *Proceedings of the CHiME-4 Workshop*, 2016, p. 4.
- [8] J. Heymann, L. Drude, and R. Haeb-Umbach, "Neural network based spectral mask estimation for acoustic beamforming," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Mar. 2016, pp. 196–200.
- [9] T. Ochiai, S. Watanabe, T. Hori, J. R. Hershey, and X. Xiao, "Unified Architecture for Multichannel End-to-End Speech Recognition With Neural Beamforming," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1274–1288, Dec. 2017.
- [10] X. Xiao, S. Watanabe, H. Erdogan, L. Lu, J. Hershey, M. L. Seltzer, G. Chen, Y. Zhang, M. Mandel, and D. Yu, "Deep beamforming networks for multi-channel speech recognition," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Mar. 2016, pp. 5745–5749.
- [11] Bo Li, Tara N. Sainath, Ron J. Weiss, Kevin W. Wilson, and Michiel Bacchiani, "Neural Network Adaptive Beamforming for Robust Multichannel Speech Recognition," in *Interspeech 2016*, Sept. 2016, pp. 1976–1980.
- [12] Z. Meng, S. Watanabe, J. R. Hershey, and H. Erdogan, "Deep long short-term memory adaptive beamforming networks for multichannel robust speech recognition," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Mar. 2017, pp. 271–275.
- [13] K. Qian, Y. Zhang, S. Chang, X. Yang, D. Florencio, and M. Hasegawa-Johnson, "Deep Learning Based Speech Beamforming," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 2018, pp. 5389–5393.
- [14] T. N. Sainath, R. J. Weiss, K. W. Wilson, A. Narayanan, and M. Bacchiani, "Factored spatial and spectral multichannel raw waveform CLDNNs," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Mar. 2016, pp. 5075–5079.
- [15] Tara N. Sainath, Arun Narayanan, Ron J. Weiss, Ehsan Vari-ani, Kevin W. Wilson, Michiel Bacchiani, and Izhak Shafran, "Reducing the Computational Complexity of Multimicrophone Acoustic Models with Integrated Feature Extraction," in *Interspeech 2016*, Sept. 2016, pp. 1971–1975.
- [16] Bo Li, Tara N. Sainath, Arun Narayanan, Joe Caroselli, Michiel Bacchiani, Ananya Misra, Izhak Shafran, Haşim Sak, Golan Pundak, Kean Chin, Khe Chai Sim, Ron J. Weiss, Kevin W. Wilson, Ehsan Vari-ani, Chanwoo Kim, Olivier Siohan, Mitchel Weintraub, Erik McDermott, Richard Rose, and Matt Shannon, "Acoustic Modeling for Google Home," in *Interspeech 2017*, Aug. 2017, pp. 399–403, ISCA.
- [17] Suyoun Kim and Ian Lane, "Recurrent Models for Auditory Attention in Multi-Microphone Distant Speech Recognition," Sept. 2016, pp. 3838–3842.
- [18] Stefan Braun, Daniel Neil, Jithendar Anumula, Enea Ceolini, and Shih-Chii Liu, "Multi-channel Attention for End-to-End Speech Recognition," in *Interspeech 2018*, Sept. 2018, pp. 17–21, ISCA.
- [19] Ehsan Vari-ani, Tara N. Sainath, Izhak Shafran, and Michiel Bacchiani, "Complex Linear Projection (CLP): A Discriminative Approach to Joint Feature Extraction and Acoustic Modeling," in *Interspeech 2016*, Sept. 2016, pp. 808–812.
- [20] Sepp Hochreiter and Jürgen Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
- [21] Duc Le, Xiaohui Zhang, Weiyi Zheng, Christian Fügen, Geoffrey Zweig, and Michael L. Seltzer, "From Senones to Chenones: Tied Context-Dependent Graphemes for Hybrid Speech Recognition," in *2019 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, Singapore, Dec. 2019.
- [22] Lukas Drude, Jahn Heymann, Christoph Boeddeker, and Reinhold Haeb-Umbach, "NARA-WPE: A Python package for weighted prediction error dereverberation in Numpy and Tensorflow for online and offline processing," in *13. ITG Fachtagung Sprachkommunikation (ITG 2018)*, Oct. 2018.
- [23] Diederik P. Kingma and Jimmy Ba, "Adam: a method for stochastic optimization," in *International Conference on Learning Representations (ICLR)*, San Diego, May 2015.
- [24] Mehryar Mohri, Fernando Pereira, and Michael Riley, "Weighted finite-state transducers in speech recognition," *Computer Speech & Language*, vol. 16, no. 1, pp. 69–88, Jan. 2002.