

# Deep CNNs Along the Time Axis With Intermap Pooling for Robustness to Spectral Variations

Hwaran Lee, Geonmin Kim, Ho-Gyeong Kim, Sang-Hoon Oh, and Soo-Young Lee

**Abstract**—Convolutional neural networks (CNNs) with convolutional and pooling operations along the frequency axis have been proposed to attain invariance to frequency shifts of features. However, this is inappropriate with regard to the fact that acoustic features vary in frequency. In this paper, we contend that convolution along the time axis is more effective. We also propose the addition of an intermap pooling (IMP) layer to deep CNNs. In this layer, filters in each group extract common but spectrally variant features, then the layer pools the feature maps of each group. As a result, the proposed IMP CNN can achieve insensitivity to spectral variations characteristic of different speakers and utterances. The effectiveness of the IMP CNN architecture is demonstrated on several LVCSR tasks. Even without speaker adaptation techniques, the architecture achieved a WER of 12.7% on the SWB part of the Hub5\*2000 evaluation test set, which is competitive with other state-of-the-art methods.

**Index Terms**—Acoustic modeling, convolutional neural networks (CNNs), intermap pooling (IMP) layer.

## I. INTRODUCTION

ACOUSTIC modeling with deep learning has demonstrated remarkable performance improvements in automatic speech recognition [1]–[4]. Deep neural networks (DNNs) are trained to label each frame of processed speech data with the state of a hidden Markov model (HMM). However, there is a difficulty due to the fact that acoustic features vary widely in frequency and articulation rate depending on harmonics of the vocal tract and characteristic speaking styles.

Efforts to effectively handle these variations can be categorized into feature-level and model-level approaches. Amongst feature-level approaches, speaker-adapted methods such as fM-LLR [5] have been proposed. Acoustic features concatenated with i-vectors, which represent speaker information, also have been employed as input for DNNs [6], [7]. Model-level approaches have employed hybrid NN-HMM systems with convolutional neural networks (CNNs) [8]–[10] and recurrent neural networks (RNNs) [11]–[13].

Manuscript received May 24, 2016; revised July 03, 2016; accepted July 08, 2016. Date of publication July 11, 2016; date of current version August 11, 2016. This work was supported by the ICT R&D program of MSIP IITP [R0126-15-1117, Core technology development of the spontaneous speech dialogue processing for the language learning]. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Daniel Povey.

H. Lee, G. Kim, H.-G. Kim, and S.-Y. Lee are with the School of Electrical Engineering, Korea Advanced Institute of Science and Technology, Daejeon 305-701, South Korea (e-mail: hwaran.lee@kaist.ac.kr; gmkim90@kaist.ac.kr; hogyong@kaist.ac.kr; sylee@kaist.ac.kr).

S.-H. Oh is with the Division of Information and Communication Convergence Engineering, Mokwon University, Daejeon 302-318, South Korea (e-mail: shoh@mokwon.ac.kr).

Color versions of one or more of the figures in this letter are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/LSP.2016.2589962

In particular, CNNs have advantages in terms of capturing local features through weight sharing while remaining robust to slight translations of these features through pooling. Structural advantages of CNNs enable the modeling of speech data without feature-level engineering, such as spectrograms or mel filterbanks. Previous researchers introduced time-delayed neural networks (TDNNs), which are CNNs with convolutions along the time axis to learn the temporal dynamics of features [14]–[16]. Other researchers have applied convolutions along the frequency axis to attain invariance to frequency-shifts [8], [17]. However, acoustic features of speech vary in frequency, so that weight sharing along the frequency axis may not be appropriate. The limited weight sharing method in which weights are convolved only within a subsection of frequency-bands has been employed in efforts to overcome this problem [10]. However, settling on appropriate band divisions and filters will require further work.

One limitation common to the preceding approaches is that most of them have employed only one or two convolution and pooling layers. Another limitation is that the relationship or topography of filters trained in supervised learning has not been intensively investigated. For unsupervised feature extraction, previous researchers imposed sparsity terms over small groups or neighborhoods in feature maps of image [18]–[20] and speech data [21], [22]. They attained topographically-organized maps of smoothly varying oriented edge filters or tonotopic disordered topography of spectrotemporal features, such as those found in the primary visual or auditory cortex (V1, A1), respectively.

In this paper, we argue that convolution along the time axis is more effective than along the frequency axis for acoustic models. In order that the network learns temporal dynamics adequately, we increase the depth of convolution layers that have small filters. Instead of frequency-axis convolution and pooling, we propose the addition of a convolutional maxout layer, namely an *intermap pooling* (IMP) layer in order to increase robustness to spectral variations. Previously, a convolutional maxout network has been proposed [23], however, it applied convolutions along the frequency axis. We show that the IMP CNNs with the time convolution reduce the word error rates more. As a result, the IMP CNNs can both model temporal dynamics and remain robust to spectral variations.

## II. CONVOLUTION NEURAL NETWORKS (CNNs)

CNNs consist of the alternation of convolution and pooling layers, and fully connected layers in the top-most layer. Let  $\mathbf{H}^{(l)}$  stand for input to the  $l$ th convolution layer having  $K$  filters, with the  $k$ th convolution filter denoted  $\mathbf{W}_k^{(l)} \in \mathbb{R}^{M \times N \times G}$  with  $M$  and  $N$  denoting a filters height and width, respectively, and  $G$  designating the number of feature maps of the input. A bias

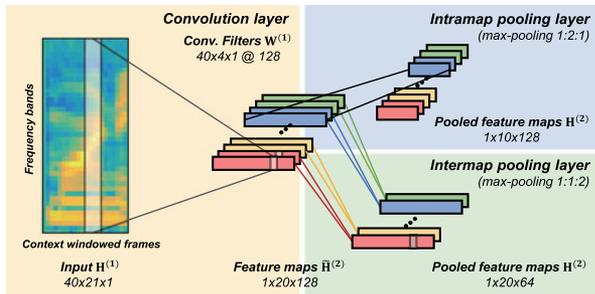


Fig. 1. Illustration of a convolution layer followed by an intramap pooling layer or an IMP layer. The sizes of convolution filters and feature maps are denoted as ‘(the freq. axis) $\times$ (the time axis) $\times$ (the feature axis)’. The number of filters is denoted after ‘@’. The pooling size is denoted as ‘(the freq. axis):(the time axis):(the feature axis)’. The convolution input is padded with zeros at both ends in the time axis in order to preserve frame length after convolutions.

term  $b_k^{(l)}$  is shared inside the  $k$ th feature map. Thus, from any input  $\mathbf{H}^{(l)}$ , the output  $\tilde{\mathbf{H}}^{(l+1)} \in \mathbb{R}^{I \times J \times K}$  can be calculated as

$$\tilde{\mathbf{H}}_{(i,j,k)}^{(l+1)} = f \left( (\mathbf{H}^{(l)} * \mathbf{W}_k^{(l)})_{(i,j)} + b_k^{(l)} \right) \quad (\text{for } 1 \leq i \leq I, 1 \leq j \leq J) \quad (1)$$

where  $I$  and  $J$  are height and width of each output feature map, and  $f$  is a nonlinear function such as sigmoid or rectified linear unit (ReLU).

An intramap pooling layer, typically called “max-pooling,” propagates the maximum value from each subregion in each feature map. For nonoverlapping subregions with height  $p$  and width  $q$ , the output from this pooling layer is given by

$$\mathbf{H}_{(i,j,k)}^{(l)} = \max_{\substack{\alpha=-p+1, \dots, 0 \\ \beta=-q+1, \dots, 0}} \tilde{\mathbf{H}}_{(i+p+\alpha, j+q+\beta, k)}^{(l)}. \quad (2)$$

Intramap pooling layers have blurring effects on feature maps, with the result being that the CNN is more robust to locally translated features.

### III. INTERMAP POOLING (IMP) LAYERS

There are several categories of acoustic features such as harmonics, formants, and on/offsets (i.e., start and end points of speech). Spectral variations of acoustic features appear as shifts in the frequency axis over time (spectrotemporal modulation). In order to ensure the robustness of our model to spectral variations, we propose the addition of a convolutional maxout layer, the IMP layer. Like the maxout networks [24], this layer *groups* the filters, and *pools* the feature maps inside a group.

Specifically, an IMP layer partitions feature maps into a set of groups. Then, each group propagates the maximum activation value at each position. Formally, the output of the  $k$ th group consisting of  $r$  consecutive feature maps is given by

$$\mathbf{H}_{(i,j,k)}^{(l)} = \max_{\gamma=-r+1, \dots, 0} \tilde{\mathbf{H}}_{(i,j,kr+\gamma)}^{(l)}. \quad (3)$$

The structural comparison of intermap and intramap pooling layers is shown in Fig. 1. Note that the method pursued in this paper does not introduce any additional learning terms except for the intermap grouping of filters.

The central idea to the IMP CNN is that the filters in each group learn common but spectrally variant features, such as frequency-shifted harmonics, and the pooled feature map is

invariant to those feature variations within the group. The pooled feature maps  $\mathbf{H}^{(l)}$  are representative of the feature maps in each group. Through supervised learning, the pooled feature maps become discriminative of features for recognizing phonemes. For a phoneme, since spectral variations among different speakers and utterances are not discriminative information, the individual filters in a group spontaneously represent common but spectrally variant features, even though the layer does not ensure this.

## IV. DEEP CNN ARCHITECTURE FOR ACOUSTIC MODELING

Since short-term temporal dynamics are shared within every frame of a given speech sample, sharing filters along the time axis is reasonable. However, sharing filters along the frequency axis may not be suitable, because features within lower frequency-band regions are significantly different from those in the higher regions. Instead of convolution along the frequency axis, our architecture employs an IMP layer following the first convolution layer. This approach demonstrates robustness not only to frequency-shifted features but also to spectrotemporally distorted features. Moreover, it does not require engineered efforts to consider the varying characteristics of different frequency-bands.

A sufficiently *deep* depth of convolution and pooling layers is necessary to precisely represent complex acoustic features with temporal and spectral variations. Individual frames also should be labeled as minutely as the number of HMM states, which is more than thousands in tri-phone modeling. However, context windowed inputs are too tiny (e.g., 21 frames) and stacking multiple intramap pooling layers decreases the feature map size in proportion to the pooling size, thereby restricting the depth of CNNs. Since previous researchers have chosen large convolution filter and intramap pooling sizes, sufficient increases in depths of CNNs have not been realized.

As illustrated in Fig. 1, the IMP CNN architecture applies convolution and intramap pooling layers only along the time axis. The pooling size of the intramap pooling layers is small so that it does not decrease temporal resolution much. Furthermore, motivated by the performance of very deep CNNs [25], we inserted convolution layers with small filters (of size  $1 \times 3$ ) between two intramap pooling layers. The combination of filters before pooling layers increases nonlinearity, and this results in a network that has rich feature expressions.

## V. EXPERIMENTAL RESULTS

### A. Experiments Setup

We conducted experiments using the 300 h Switchboard-I Release 2 (SWBD) dataset [26] which is conversational telephone speech task as well as the Wall Street Journal (WSJ) corpus [27] and Aurora4 database which are read speech. We used the 81-h training dataset (SI-284) of the WSJ corpus. The Aurora4 database is a subset of the WSJ in which clean utterances are added with different noise types and/or convolved with microphone distortions. The following results are for the trained IMP CNN on the multiconditioned training dataset.

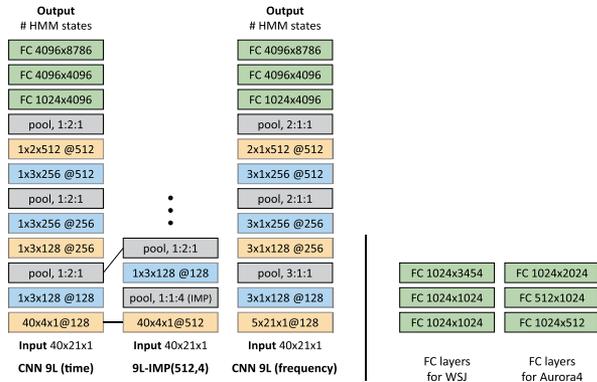


Fig. 2. Configurations of CNN architectures for SWBD. The ReLU nonlinearity function is used on the top of every activation. We made 6, 9, 12, and 15 layers by excluding or repeating the blue colored layers. For WSJ and Aurora4 datasets, the fully connected layers have smaller hidden neurons, and remainder are the same.

The raw speech signal is processed via short-time Fourier transform (STFT) with a 25 ms Hamming window and 10 ms window shifts. We used 40-dimensional (40-D) log-mel filter bank features without the energy coefficient, and concatenated frames with a context window size of 21 ( $\pm 10$  frames) to feed them into networks as inputs. We trained the GMM-HMM system over fMLLR features. The forced alignment of each frame by the GMM-HMM baseline system is the target label of the neural networks.

After random initialization of weights and biases from the Gaussian distribution  $\mathcal{N}(0, 0.01)$  and  $\mathcal{N}(0, 0.5)$ , respectively, the CNNs were optimized by the stochastic gradient descent (SGD) method. In particular, for CNNs deeper than nine-layers, we faced with infeasible training, because each layer back-propagates errors by multiplying its small initial weights, resulting in vanishing gradients. Therefore, we increased the standard deviation ( $\sigma$ ) of the Gaussian distribution in lower layers. Each layer is trained with a momentum of 0.9, an L2-decay term of 0.0005, and minibatch size of 512. After one epoch of training, the trained model is accepted if the validation cost decreases. Otherwise, the trained model is rejected and training starts again from the latest accepted model with a halved learning rate. The initial learning rate is 0.01, and the training stops after 50 epochs. Our implementation is developed upon the KALDI toolkit [28].

For the SWBD task, we decode speech using a trigram language model (LM) of 30k vocabularies which is trained on 3M words, and then we rescore the decoding results using 4-gram LM which is trained on Fisher English Part 1 transcripts [29]. For the WSJ and Aurora4 corpus, we used a 146K word extended dictionary and the trigram pruned language model which is exactly the same as the “s5” recipe in the KALDI.

### B. Convolution Axis and Depth of CNNs

Fig. 3 shows the decoding results of CNNs on SWBD evaluation sets with various depths, from 6 layers up to 15 layers (configurations are described in Fig. 2). Deeper CNNs produced lower WERs, with the 15-layer CNN achieving a WER of 12.8% for SWB and 18.6% for total evaluation sets. Moreover, it is validated that convolution along the time axis always

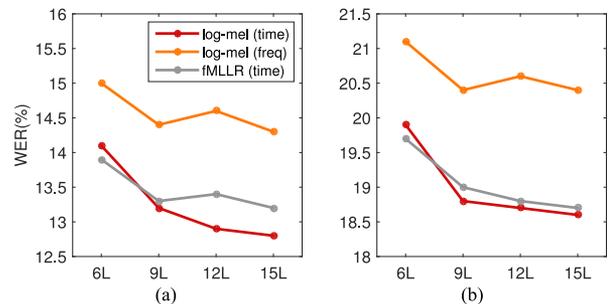


Fig. 3. Decoding results of CNNs on the Switchboard evaluation sets (Hub5\*2000). CNNs with different convolution axis (time, frequency) and input features (log-mel, fMLLR) are compared. (a) Switchboard subset (b) Total (SWB and CallHome).

TABLE I  
WERS(%) OF CNNs WITH AN IMP LAYER AND AN OVERLAPPING IMP (IMPO) ON SWBD

Network	SWB	Total
9L	13.2	18.8
9L-IMP(128,2)	13.2	18.8
9L-IMP(256,2)	13.2	18.7
9L-IMP(512,4)	<b>12.7</b>	<b>18.5</b>
9L-IMP(768,6)	12.9	18.7
9L-IMPO(512,4)	13.0	18.6

TABLE II  
IMP IN DIFFERENT CONVOLUTION AXIS

	SWB		Total	
	9L	9L-IMP	9L	9L-IMP
log-mel (time)	13.2	<u>12.7</u>	18.8	<u>18.5</u>
log-mel (freq.)	<u>14.4</u>	14.7	<u>20.4</u>	20.7
fMLLR (time)	<u>13.3</u>	13.5	<u>19.0</u>	19.1

outperforms convolution along the frequency axis. Furthermore, CNNs trained over log-mel features had lower WER as fMLLR features when CNN has more than nine layers. These results show that weight sharing along the time axis more effectively reduces the WER, and that increased nonlinearity obviates pre-processing for speaker adaptation.

### C. IMP CNNs

Decoding results of IMP-CNNs with different numbers of maps and pooling sizes are compared in Table I. We further investigated an IMP layer in which groups overlap (IMPO) each other with a stride of one. All CNNs with an IMP layer performed better than the 9-layer CNN without. Especially, the “9L-IMP(512, 4)” CNN performs the best with a WER of 12.7% for SWB test set, showing a 3.78% relative improvement over the 9-layer CNN. Note Table II that when the IMP layer is applied to CNNs along the frequency axis or over fMLLR features, performance declines. In addition, IMP CNN performed well on the WSJ and Aurora4 corpus as shown in Tables III and IV, respectively. It is remarkable that IMP layers contribute robustness to spectral variations in both clean and noisy conditions.

TABLE III  
WERS(%) OF CNNs WITH AN IMP LAYER ON WSJ

	9L	9L-IMP	Rel. (%)
Eval'92	4.27	<b>3.93</b>	7.96

TABLE IV  
WERS(%) OF CNNs WITH AN IMP LAYER ON AURORA4

	Clean	Noise	Channel	Channel+Noise	Avg.
9L	3.36	7.05	8.14	18.05	11.58
9L-IMP	<b>3.14</b>	<b>6.64</b>	<b>7.86</b>	<b>17.88</b>	<b>11.29</b>
Rel. (%)	6.55	5.82	3.44	0.92	2.50

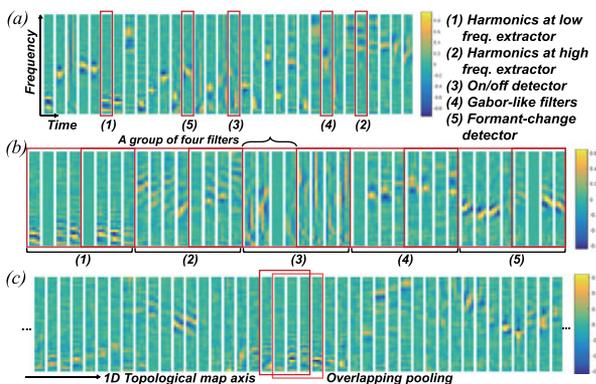


Fig. 4. Trained filters of the first convolution layer. (a) Filters of the 9-layer CNN. The top 22 filters that have the biggest L2-norm are sorted in decreasing order. (b) 10 groups of four filters in “9L-IMP(512, 4)”. The order of groups (red boxes) is arranged according to the feature category of each group. (c) A part of successive filters in “9L-IMP(512, 4)”.

#### D. Analysis on Learnt Filters

Learnt filters of the first convolution layer are visualized in Fig. 4(a). There are *five categories of spectrotemporal features* in the filters. (1) Harmonic features are narrow in the low frequency-region and (2) broad in the high-frequency region. (3) The on/off-set detecting filters are temporally selective, but are also sensitive to several frequencies. (4) The features of Gabor-like filters are centered on some frequency-bands, which presumably detect formants (5). The features of formant changes are directional diagonal lines, spectrotemporal modulations, in the middle frequency-bands. Note that different features appear in different frequency-bands, and that local features of each type have different bandwidth sizes.

The trained filters in each group of the IMP layer are presented in Fig. 4(b). Importantly, most filters in a group belong to a common category. For example, the filters in harmonic extractor and formant change detector groups have marginally shifted features on the frequency axis. This figure verifies that the IMP lead the filters of a group to extract common but spectrally variant features, although there are no additional architectural constraints to guarantee this.

The consecutive trained filters of the IMPO layer are drawn in Fig. 4(c). The filters form a 1-D topological map, where neighboring filters respond to similar spectrotemporal features.

TABLE V  
COMPARISON OF WERS(%) FOR DIFFERENT MODELS ON SWBD

Input features	Model	SWB	Total
fMLLR	GMM-HMM	19.5	26.7
	Maxout 7L	14.2	20.0
	log-mel filterbanks	14.6	20.7
	CNN 9L	13.2	18.8
	CNN 9L-IMP(512, 4)	<u>12.7</u>	<b>18.5</b>
	CNN 15L	12.8	18.6
MFCC + i-vectors	TDNN 4L [16]	12.9	19.2
VTL-warped log-mel	CNN 8L (2conv+6fc) [34]	<b>12.6</b>	–
	CNN 13L (10conv+3fc) [35]	<b>11.8</b>	–

Along the topological map axis, filters appear discontinuously between feature categories, reflecting the fact that feature categories become definitely distinguishable to the system as it is trained. In recent neurophysiological studies, there is consensus that multiple tonotopic maps exist in the human auditory system [30]. However, a few studies suggest that this topography includes other sound features, such as temporal, spectral, and joint modulations [31]–[33]. The trained topography may provide a clue as to how human auditory neurons organize to efficiently process information in A1.

#### E. Comparison of the IMP CNN

For comparison, we trained max-out networks that have seven layers with 2000 hidden neurons and 400 groups on both fMLLR and filter-bank features. The comparison of the decoding results is summarized in Table V. The “9L-IMP(512, 4)” IMP CNN improved on the GMM-HMM baseline (19.5%) and the max-out network (14.6%), demonstrating a 34.87% and 13.01% relative improvement, respectively. Also, it performs on par with a 15-layer CNN, i.e., a nonIMP CNN with six additional convolution layers. Finally, the IMP CNN is compared with the TDNN [16] and the CNNs which employed 2-D convolutions [34], [35]. Note that we only compare other previous results without any sequence training such as sMBR. Even though our deep CNN did not use any speaker adaptation techniques, it yielded a comparative word error rate simply by employing IMP and by increasing depths.

## VI. CONCLUSION

In this paper, the present experiments demonstrate that convolution along the time axis is more effective than along the frequency axis when processing speech. Depth in convolution layers is crucial for the sufficient representation of the complex temporal dynamics inherent in the acoustic features of speech. In order to achieve greater robustness to spectral variations in speech recognition, we proposed the addition of IMP to CNNs. Through visualization of the trained filters, we verified that filters grouped together learn similar spectrotemporal features and form a topological map. In the end, even without any speaker adaptation techniques, the proposed IMP CNN delivered competitive performance on the Switchboard, WSJ, and Aurora4 databases.

## REFERENCES

- [1] A.-R. Mohamed, G. E. Dahl, and G. Hinton, "Acoustic modeling using deep belief networks," *IEEE Trans. Audio, Speech, Language Process.*, vol. 20, no. 1, pp. 14–22, Jan. 2012.
- [2] G. Dahl, D. Yu, L. Deng, and A. Acero, "Large vocabulary continuous speech recognition with context-dependent DBN-HMMs," *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2011, pp. 4688–4691.
- [3] J. Lee and S.-Y. Lee, "Deep learning of speech features for improved phonetic recognition," *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2011, pp. 1249–1252.
- [4] G. Hinton *et al.*, "Deep neural networks for acoustic modeling in speech recognition," *Signal Process. Mag.*, vol. 29, no. 6, pp. 82–97, 2012.
- [5] S. P. Rath, D. Povey, K. Vesel, and J. H. Cernock, "Improved feature processing for deep neural networks," in *Proc. 14th Annu. Conf. Int. Speech Commun. Assoc.*, 2013, pp. 1–5.
- [6] G. Saon, H. Soltan, D. Nahamoo, and M. Picheny, "Speaker adaptation of neural network acoustic models using i-vectors," *Proc. IEEE Work. Autom. Speech Recogn. Understanding*, 2013, pp. 55–59.
- [7] A. Senior and I. Lopez-Moreno, "Improving DNN speaker independence with I-vector inputs," *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2014, pp. 225–229.
- [8] O. Abdel-Hamid, A.-R. Mohamed, H. Jiang, and G. Penn, "Applying convolutional neural networks concepts to hybrid NN-HMM model for speech recognition," *IEEE Int. Conf. Acoust. Speech Signal Process.*, 2012, pp. 4277–4280.
- [9] O. Abdel-Hamid, L. Deng, and D. Yu, "Exploring convolutional neural network structures and optimization techniques for speech recognition," *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2013, pp. 3366–3370.
- [10] O. Abdel-Hamid, A.-R. Mohamed, H. Jiang, L. Deng, G. Penn, and D. Yu, "Convolutional neural networks for speech recognition," *IEEE/ACM Trans. Audio, Speech, Language Process.*, vol. 22, no. 10, pp. 1533–1545, Oct. 2014.
- [11] A. Graves, A.-R. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2013, pp. 6645–6649.
- [12] A. Graves, N. Jaitly, and A. R. Mohamed, "Hybrid speech recognition with deep bidirectional LSTM," *Proc. IEEE Workshop Autom. Speech Recogn. Understanding*, 2013, pp. 273–278.
- [13] F. Beaufays, H. Sak, and A. Senior, "Long short-term memory recurrent neural network architectures for large scale acoustic modeling has," *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2014, pp. 338–342.
- [14] A. Waibel, T. Hanazawa, G. E. Hinton, K. Shikano, and K. J. Lang, "Phoneme recognition using time-delay neural networks," *IEEE Trans. Acoust. Speech Signal Process.*, vol. 37, no. 3, pp. 328–339, Mar. 1989.
- [15] H. Lee, P. Pham, Y. Largman, and A. Ng, "Unsupervised feature learning for audio classification using convolutional deep belief networks," *Proc. Adv. Neural Inf. Process. Syst.*, 2009, pp. 1–9.
- [16] V. Peddinti, D. Povey, and S. Khudanpur, "A time delay neural network architecture for efficient modeling of long temporal contexts," *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2015, pp. 2–6.
- [17] T. N. Sainath, A.-R. Mohamed, B. Kingsbury, and B. Ramabhadran, "Deep Convolutional neural networks for LVCSR," *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2013, pp. 10–14.
- [18] A. Hyvärinen and H. Patrik, "Emergence of phase and shift invariant features by decomposition of natural images into independent feature subspaces," *Neural Comput.*, vol. 12, no. 7, pp. 1705–1720, 2000.
- [19] A. Hyvärinen, P. O. Hoyer, and M. Inki, "Topographic independent component analysis," *Neural Comput.*, vol. 13, no. 7, pp. 1527–1558, 2001.
- [20] K. Kavukcuoglu, M. Ranzato, R. Fergus, and Y. Le-Cun, "Learning invariant features through topographic filter maps," *Comput. Vis. Pattern Recog.*, 2009, pp. 1605–1612.
- [21] T. Kim and S.-Y. Lee, "Learning self-organized topology-preserving complex speech features at primary auditory cortex," *Neurocomputing*, vol. 65–66, pp. 793–800, 2005.
- [22] H. Terashima and M. Okada, "The topographic unsupervised learning of natural sounds in the auditory cortex," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 1–9.
- [23] M. Cai, Y. Shi, J. Kang, J. Liu, and T. Su, "Convolutional maxout neural networks for low-resource speech recognition," *Proc. 9th Int. Symp. Chinese Spok. Language Process.*, 2014, pp. 133–137.
- [24] I. J. Goodfellow, D. Warde-Farley, M. Mirza, A. Courville, and Y. Bengio, "Maxout networks," in *Proc. 30th Int. Conf. Mach. Learn.*, 2013, pp. 1319–1327.
- [25] K. Simonyan, and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *ICLR*, 2015.
- [26] J. J. Godfrey, E. C. Holliman, and J. McDaniel, "SWITCHBOARD telephone speech corpus for research and development," in *Proc. IEEE Int. Conf. Acoust. Speech, Signal Process.*, 1992, pp. 517–520.
- [27] D. B. Paul and J. M. Baker, "The design for the wall street journal-based CSR corpus," in *Proc. Workshop Speech Natural Language*, 1994, pp. 357–362.
- [28] D. Povey *et al.*, "The kaldi speech recognition toolkit," *Proc. IEEE Workshop IEEE Autom. Speech Recogn. Understanding*, 2011, pp. 1–4.
- [29] C. Cieri, D. Miller, and K. Walker, "The fisher corpus: A resource for the next generations of speech-to-text," *Proc. 4th Int. Conf. Language Resources Eval.*, 2004, pp. 69–71.
- [30] M. Saenz and D. R. Langers, "Tonotopic mapping of human auditory cortex," *Hearing Res.*, vol. 307, pp. 42–52, 2014.
- [31] B. Barton, J. H. Venezia, K. Saberi, G. Hickok, and A. A. Brewer, "Orthogonal acoustic dimensions define auditory field maps in human cortex," *Proc. Natl. Acad. Sci. USA*, vol. 109, pp. 20738–20743, 2012.
- [32] M. Herdener *et al.*, "Spatial representations of temporal and spectral sound cues in human auditory cortex," *Cortex*, vol. 49, no. 10, pp. 2822–2833, 2013.
- [33] R. Santoro *et al.*, "Encoding of natural sounds at multiple spectral and temporal resolutions in the human auditory cortex," *PLoS Comput. Biol.*, vol. 10, no. 1, 2014, Art. no. e1003412.
- [34] G. Saon, H.-K. J. Kuo, S. Rennie, and M. Picheny, "The IBM 2015 english conversational telephone speech recognition system," in *Proc. 16th Annu. Conf. Int. Speech Commun. Assoc.*, 2015.
- [35] T. Sercu, C. Puhresch, B. Kingsbury, and Y. LeCun, "Very deep multilingual convolutional neural networks for LVCSR," in *Proc. Neural Netw. Trends Speech Recog.*, 2016, pp. 2–6.