

Recent Progresses in Deep Learning Based Acoustic Models

Dong Yu and Jinyu Li

Abstract—In this paper, we summarize recent progresses made in deep learning based acoustic models and the motivation and insights behind the surveyed techniques. We first discuss models such as recurrent neural networks (RNNs) and convolutional neural networks (CNNs) that can effectively exploit variable-length contextual information, and their various combination with other models. We then describe models that are optimized end-to-end and emphasize on feature representations learned jointly with the rest of the system, the connectionist temporal classification (CTC) criterion, and the attention-based sequence-to-sequence translation model. We further illustrate robustness issues in speech recognition systems, and discuss acoustic model adaptation, speech enhancement and separation, and robust training strategies. We also cover modeling techniques that lead to more efficient decoding and discuss possible future directions in acoustic model research.

Index Terms—Attention model, convolutional neural network (CNN), connectionist temporal classification (CTC), deep learning (DL), long short-term memory (LSTM), permutation invariant training, speech adaptation, speech processing, speech recognition, speech separation.

I. INTRODUCTION

IN the past several years, there has been significant progress in automatic speech recognition (ASR) [1]–[21]. These progresses have led to ASR systems that surpassed the threshold for adoption in many real-world scenarios and enabled services such as Google Now, Microsoft Cortana, and Amazon Alexa. Many of these achievements are powered by deep learning (DL) techniques. Readers are referred to Yu and Deng 2014 [22] for a comprehensive summary and detailed description of the technology advancements in ASR made before 2015.

In this paper, we survey new developments happened in the past two years with an emphasis on acoustic models. We discuss motivations and core ideas of each interesting work surveyed. More specifically, in Section II we illustrate improved DL/HMM (hidden Markov model) hybrid acoustic models that employ deep recurrent neural networks (RNNs) and deep convolutional neural networks (CNNs). These hybrid

models can better exploit contextual information than feed-forward deep neural networks (DNNs) and thus lead to new state-of-the-art recognition accuracy. In Section III we describe acoustic models that are designed and optimized end-to-end with no or less non-learn-able components. We first discuss the models in which audio waveforms are directly used as the input feature so that the feature representation layer is automatically learned instead of manually designed. We then depict models that are optimized using the connectionist temporal classification (CTC) criterion which allows a sequence-to-sequence direct mapping. Following that we analyze sequence-to-sequence translation models that are built upon the attention mechanism. We devote Section IV to discuss techniques that can improve robustness with focuses on adaptation techniques, speech enhancement and separation techniques, and robust training techniques. In Section V we describe acoustic models that support efficient decoding and cover frame-skipping and model compression through teacher-student training and quantization. We propose core problems to work on and potential future directions in solving them in Section VI.

II. ACOUSTIC MODELS EXPLOITING VARIABLE-LENGTH CONTEXTUAL INFORMATION

The DL/HMM hybrid model [1]–[5] is the first deep learning architecture that succeeded in ASR and is still the dominant model used in industry. Several years ago, most hybrid systems are DNN based. As reported in [3], one of the important factors that lead to superior performance in the DNN/HMM hybrid system is its ability to exploit contextual information. In most systems, a window of 9 to 13 frames (left/right context of 4–6 frames) of features are used as the input to the DNN system to exploit the information from neighboring frames to improve the accuracy.

However, the optimal length of contextual information may vary for different phones and speaking speed. This indicates that using fixed-length context window, as in the DNN/HMM hybrid system, may not be the best choice to exploit contextual information. In recent years people have proposed new models that can exploit variable-length contextual information more effectively. The most important two models use deep RNNs and CNNs.

A. Recurrent Neural Networks

Feed-forward DNNs only consider information in a fixed-length sliding window of frames and thus cannot exploit long-range correlations in the speech signal. On the other hand, RNNs can encode sequence history in their internal states,

Manuscript received April 11, 2017; accepted May 24, 2017. Recommended by Associate Editor Qinglai Wei. (Corresponding author: Dong Yu.)

Citation: D. Yu and J. Li, "Recent progresses in deep learning based acoustic models," *IEEE/CAA J. of Autom. Sinica*, vol. 4, no. 3, pp. 396–409, Jul. 2017.

D. Yu is with the Tencent AI Lab, Tencent, Bellevue, WA 98034, USA (e-mail: dongyu@ieee.org).

J. Li is with the Microsoft AI and Research, Microsoft, Redmond, WA 98052, USA (e-mail: jinyuli@microsoft.com).

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

Digital Object Identifier 10.1109/JAS.2017.7510508

and thus have the potential to predict phonemes based on all the speech features observed up to the current frame. Unfortunately, simple RNNs, depending on the largest eigenvalue of the state-update matrix, may have gradients which either increase or decrease exponentially over time. Hence, the basic RNNs are difficult to train, and in practice can only model short-range effects.

Long short-term memory (LSTM) RNNs [23] were developed to overcome these problems. LSTM-RNNs use input, output and forget gates to control information flow so that gradients can be propagated in a stable fashion over relatively longer span of time. These networks have been shown to outperform DNNs on a variety of ASR tasks [8], [24]–[27]. Note that there is another popular RNN model, called gated recurrent unit (GRU), which is simpler than LSTM but is also able to model the long short-term correlation. Although GRU has been shown effective in several machine learning tasks [28], it is not widely used in ASR tasks.

At the time step t , the vector formulas of the computation of LSTM units can be described as:

$$\mathbf{i}_t = \sigma(\mathbf{W}_{ix}\mathbf{x}_t + \mathbf{W}_{ih}\mathbf{h}_{t-1} + \mathbf{p}_i \odot \mathbf{c}_{t-1} + \mathbf{b}_i) \quad (1a)$$

$$\mathbf{f}_t = \sigma(\mathbf{W}_{fx}\mathbf{x}_t + \mathbf{W}_{fh}\mathbf{h}_{t-1} + \mathbf{p}_f \odot \mathbf{c}_{t-1} + \mathbf{b}_f) \quad (1b)$$

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \phi(\mathbf{W}_{cx}\mathbf{x}_t + \mathbf{W}_{ch}\mathbf{h}_{t-1} + \mathbf{b}_c) \quad (1c)$$

$$\mathbf{o}_t = \sigma(\mathbf{W}_{ox}\mathbf{x}_t + \mathbf{W}_{oh}\mathbf{h}_{t-1} + \mathbf{p}_o \odot \mathbf{c}_t + \mathbf{b}_o) \quad (1d)$$

$$\mathbf{h}_t = \mathbf{o}_t \odot \phi(\mathbf{c}_t) \quad (1e)$$

where \mathbf{x}_t is the input vector. The vectors \mathbf{i}_t , \mathbf{o}_t , \mathbf{f}_t are the activation of the input, output, and forget gates, respectively. The $\mathbf{W}_{\cdot x}$ and $\mathbf{W}_{\cdot h}$ terms are the weight matrices for the inputs \mathbf{x}_t and the recurrent inputs \mathbf{h}_{t-1} , respectively. The \mathbf{p}_i , \mathbf{p}_o , \mathbf{p}_f are parameter vectors associated with peephole connections. The functions σ and ϕ are the logistic sigmoid and hyperbolic tangent nonlinearity, respectively. The operation \odot represents element-wise multiplication of vectors.

It is popular to stack LSTM layers to get better modeling power [8]. However, an LSTM-RNN with too many vanilla LSTM layers is very hard to train and there still exists the gradient vanishing issue if the network goes too deep. This issue can be solved by using either highway LSTM or residual LSTM.

In the highway LSTM [29], memory cells of adjacent layers are connected by gated direct links which provide a path for information to flow between layers more directly without decay. Therefore, it alleviates the gradient vanishing issue and enables the training of much deeper LSTM-RNN networks.

Residual LSTM [30], [31] uses shortcut connections between LSTM layers, and hence also provides a way to alleviating the gradient vanishing problem. Different from highway LSTM which uses gates to guide the information flow, residual LSTM is more straightforward with the direct shortcut path, similar to Residual CNN [32] which recently achieves great success in the image classification task.

Typically, log Mel-filter-bank features are often used as the input to the neural-network-based acoustic models [33], [34]. Switching two filter-bank bins will not affect the performance of the DNN or LSTM. However, this is not the case when a human reads a spectrogram: a human relies on both patterns

that evolve over time and frequency to predict phonemes. This inspired the proposal of a 2-D, time-frequency (TF) LSTM [35], [36] which jointly scans the speech input over the time and frequency axes to model spectro-temporal warping, and then uses the output activation as the input to the traditional time LSTM. The joint time-frequency modeling provides better normalized features for the upper layer time LSTMs. This has been verified effective and robust to distortion at both Microsoft and Google on large-scale tasks [35]–[37].

Highway LSTM has gates on both the temporal and spatial directions while TF LSTM has gates on both the temporal and spectral directions. It is desirable to have a general LSTM structure that works along all directions. Grid LSTM [38] is such a general LSTM which arranges the LSTM memory cells into a multidimensional grid. It can be considered as a unified way of using LSTM for temporal, spectral, and spatial computation. Grid LSTM has been studied for temporal and spatial computation in [39] and temporal and spectral computation in [37].

Although bi-directional LSTMs (BLSTMs) perform better than uni-directional LSTMs by using the past and future context information [8], [40], they are not suitable for real-time systems since the recognition can happen only after the whole utterance has been observed. For this reason, models, such as latency-controlled BLSTM (LC-BLSTM) [29] and row-convolution BLSTM (RC-BLSTM), that bridge between uni-directional LSTMs and BLSTMs have been proposed. In these models, the forward LSTM is still kept as is. However, the backward LSTM is replaced by either a backward LSTM with at most N -frames of lookahead as in the LC-BLSTM case, or a row-convolution operation that integrates information in the N -frames of lookahead. By carefully choosing N we can balance between recognition accuracy and latency. Recently, LC-BLSTM was improved in [41] to speed up the evaluation and to enable real-time online speech recognition by using better network topology to initialize the BLSTM memory cell states.

B. Convolutional Neural Networks

Another model that can effectively exploit variable-length contextual information is the convolutional neural network (CNN) [42], in the center of which is the convolution operation (or layer). The input to the convolution operation is usually a three-dimensional tensor (row, column, channel) for speech recognition but can be lower or higher dimensional tensors for other applications. Each channel of the input and output of the convolution operation can be considered as a view of the same data. In most setups, all channels have the same size (height, width).

The filters in the convolution operation are called kernels, which are four-dimensional tensors (kernel height, kernel width, input channel, output channel) in our case. There are in total $C_x \times C_v$ kernels, where C_x is the number of input channels and C_v is the number of output channels. The kernels are applied to local regions called receptive fields in an input

image along all channels. The value after the convolution operation is

$$v_{ij\ell}(\mathbf{K}, \mathbf{X}) = \sum_n \text{vec}(\mathbf{K}_{n\ell}) \cdot \text{vec}(\mathbf{X}_{ijn}) \quad (2)$$

for each output channel ℓ and input slice (i, j) (the i th step along the vertical direction and j th step along the horizontal direction), where $\mathbf{K}_{n\ell}$ of size (H_k, W_k) is a kernel matrix associated with input channel n and output channel ℓ and has the same size as the input image patch \mathbf{X}_{ijn} of channel n , $\text{vec}(\cdot)$ is the vector formed by stacking all the columns of the matrix, and \cdot is the inner product of two vectors. It is obvious that each output pixel is a weighted sum of all pixels across all channels in an input patch. Since each input pixel can be considered as a weak pattern detector, each output pixel is just a boosted detector exploiting all information in the input patch.

The kernel is shared across all input patches and moves along the input image with strides S_r and S_c at the vertical and horizontal direction, respectively. When the strides are larger than 1, the convolution operation also subsamples, in addition to convolving, the input image and leads to a lower-resolution image that is less sensitive to the small pattern shift inside the input patch. The translational invariance can be further improved when some kind of aggregation operations are applied after the convolution operation. Typical aggregation operations are max-pooling and average-pooling. The aggregation operations often go with subsampling to reduce resolution. Due to the built-in translational invariability CNNs can exploit variable-length contextual information along both frequency and time axes. It is obvious that if only one convolution layer is used, the translational variability that the system can tolerate is limited. To allow for more powerful exploitation of the variable-length contextual information, convolution operations (or layers) can be stacked.

The time delay neural network (TDNN) [43] was the first model that exploits multiple CNN layers for ASR. In this model, convolution operations are applied to both time and frequency axes. However, the early TDNNs are neural-network-only solutions that do not integrate with HMMs and are hard to be used in large vocabulary continuous speech recognition (LVCSR).

After the successful application of DNNs to LVCSR, CNNs were reintroduced under the DL/HMM hybrid model architecture [5], [7], [11], [14], [17], [44]–[46]. Because HMMs in the hybrid model already have strong ability to handle variable-length utterance problem in ASR, CNNs were reintroduced initially to deal with variability at the frequency axis only [5], [7], [44], [45]. The goal was to improve robustness against vocal tract length variability between different speakers. Only one to two CNN layers were used in these early models, stacked with additional fully-connected DNN layers. These models have shown around 5% relative recognition error rate reduction compared to the DNN/HMM systems [7]. Later, additional RNN layers, e.g., LSTMs, were integrated into the model to form so called CNN-LSTM-DNN (CLDNN) [10] and CNN-DNN-LSTM (CDL) architectures. The RNNs in these models can help to exploit the variable-length contextual

information since CNNs in these models only deal with frequency-axis variability. CLDNN and CDL both achieved additional accuracy improvement over CNN-DNN models.

Researchers quickly realized that dealing with variable-length utterance is different from exploiting variable-length contextual information. TDNNs, which convolve along both the frequency and time axes and thus exploit variable-length contextual information, attracted new attentions, this time under the DL/HMM hybrid architecture [13], [47] and with variations such as row convolution [15] and feedforward sequential memory network (FSMN) [16]. Similar to the original TDNNs, these models stack several CNN layers along the frequency and time-axis, with a focus on the time-axis, to account for speaking rate variation. But unlike the original TDNNs, the TDNN/HMM hybrid systems can recognize large vocabulary continuous speech very effectively.

More recently, primarily motivated by the successes in image recognition, various architectures of deep CNNs [14], [17], [46], [48] have been proposed and evaluated for ASR. The premise is that spectrograms can be seen as images with special patterns from which experienced people can tell what has been said. In deep CNNs, each higher layer is a weighted sum of nonlinear transformation of a window of lower layers and thus covers longer contexts and operates on more abstract patterns. Lower CNN layers capture local simple patterns while higher CNN layers detect broader, abstract, and more complicated patterns. Smaller kernels combined with more layers allow deep CNNs to exploit longer-range dependency information along both time and frequency axes more effectively. Empirically deep CNNs are compatible to BLSTMs [19], which in turn outperform unidirectional LSTMs. However, unlike BLSTMs which suffer from long latency, deep CNNs have limited latency and are better suited for real-time systems if the computation cost can be controlled.

Training and evaluation of deep CNNs is very time consuming, esp. if we treat each window of frames independently, under which condition there are significant duplication of computations. To speedup the computation we can treat the whole utterance as a single input image and thus reuse the intermediate computation results. Even better, if the deep CNN is designed so that the stride at each layer is long enough to cover the whole kernel, similar to the CNNs with layer-wise context expansion and attention (LACE) [17]. Such model, called dilated CNN [46], allows to exploit longer-range information with less number of layers and can significantly reduce the computational cost. Dilated CNN has outperformed other deep CNN models on the switchboard task [46].

Note that deep CNNs can be used together with RNNs and under frameworks such as connectionist temporal classification (CTC) that we will discuss in Section III-B.

III. ACOUSTIC MODELS WITH END-TO-END OPTIMIZATION

The models discussed in the previous section are DNN/HMM hybrid models in which the two components DNN and HMM are usually optimized separately. However, speech recognition is a sequential recognition problem. It is

not surprising that better recognition accuracy may be achieved if all components in a model are jointly optimized. It is even better if the model can remove all manually designed components such as basic feature representation and lexicon design.

A. Automatically Learned Audio Feature Representation

It is always arguable that the manually-designed log Mel-filter-bank feature is optimal for speech recognition. Inspired by the end-to-end processing in the machine learning community, there are always efforts [49]–[52] trying to replace the Mel-filter-bank extraction by directly learning filters using a network to process the raw speech waveforms and training it jointly with the recognizer network. Among these efforts, the CLDNN [10] on raw waveform [52] seems to be more promising as it got slight gain over the log Mel-filter-bank feature while the other works did not. More importantly, it serves a good foundation of the multichannel processing with raw waveforms.

The most critical thing for raw waveform processing is using a representation that is invariant to small phase shift because the raw waveforms are perceptually identical if the only difference is a small phase shift. To achieve the phase invariance, a time convolutional layer is applied to the raw waveform and then pooling is done over the entire time length of the time-convolved output signal. This process reduces the temporal variation (hence phase invariant) and is very similar to the Gammatone filterbank extraction. The pooled outputs can be considered as the filter-bank outputs, on which the standard CLDNN [10] is applied. Interestingly, the same idea was recently applied to anti-spoofing speaker verification task with significant gain [53].

With the application of deep learning models, now the ASR systems on close-talking scenario perform very well. The research interest is shifted to the far-field ASR which needs to handle both additive noise and reverberation. The current dominant approach is still using the traditional beamforming method to process the waveforms from multiple microphones, and then inputting the beamformed signal into acoustic models [54]. Efforts have also been made to use deep learning models to perform beamforming and jointly train the beamforming and recognizer networks [55]–[58]. In [55], the beamforming network and recognizer networks were trained in a sequence by first training the beamforming network, then training the recognizer network with the beamformed signal, and finally jointly training both networks.

In [56]–[58], both networks were jointly trained in a more end-to-end fashion by extending the aforementioned CLDNN to raw waveform. In the first layer, multiple time convolution filters are used to map the raw waveforms from multiple microphones into a single time-frequency representation [56]. Then the output is passed to the upper layer CLDNN for phoneme classification. Later, the joint network is improved by factorizing the spatial and spectral selectivity of bottom layer network into a spatial filtering layer and a spectral filtering layer. The factored network brings accuracy improvement with additional computational cost, which later was reduced by

converting the time-domain convolution into the frequency-domain product [59].

B. Connectionist Temporal Classification

Speech recognition task is a sequence-to-sequence translation task, which maps the input waveform to a final word sequence or an intermediate phoneme sequence. What the acoustic model really should care is the output word or phoneme sequence, instead of the frame-by-frame labeling that is optimized in the traditional cross entropy (CE) training. The connectionist temporal classification (CTC) approach [9], [60], [61] was introduced to adopt this view and maps the speech input frames into an output label sequence. To deal with the issue that the number of output labels is smaller than that of input speech frames in speech recognition tasks, CTC introduces a special blank label and allows for repetition of labels to force the output and input sequences to have the same length.

Denote \mathbf{x} as the speech input sequence, \mathbf{l} as the original label sequence, and $B^{-1}(\mathbf{l})$ represents all the CTC paths mapped from \mathbf{l} . The CTC loss function is defined as the sum of negative log probabilities of correct labels as

$$L_{\text{CTC}} = -\ln P(\mathbf{l}|\mathbf{x}) \quad (3)$$

where

$$P(\mathbf{l}|\mathbf{x}) = \sum_{\mathbf{z} \in B^{-1}(\mathbf{l})} P(\mathbf{z}|\mathbf{x}). \quad (4)$$

With the conditional independence assumption, $P(\mathbf{z}|\mathbf{x})$ can be decomposed into a product of posterior from each frame as

$$P(\mathbf{z}|\mathbf{x}) = \prod_{t=1}^T P(z_t|\mathbf{x}). \quad (5)$$

The calculation of $P(z_t|\mathbf{x})$ is via the forward-backward process in [62].

In [60], CTC with context-dependent phone output units have been shown to outperform CTC with monophone [9], and to perform in par with the LSTM model with cross entropy criterion when training data is large enough. One attractive characteristics of CTC is that we can choose output units such as syllables and words that are larger than phoneme. This implies that the input features can be constructed with a sampling rate that is larger than 10 ms. For example, in [60], three 10 ms frames are stacked together as the input to CTC models. By doing so, the acoustic score evaluation during decoding happens every 30 ms, 3 times faster than the traditional systems that operate on 10 ms frame shift.

CTC provides a path to end-to-end optimization of acoustic models. In the deep speech [15], [63] and EESSEN [64], [65] work, the end-to-end speech recognition systems were explored to directly predict characters instead of phonemes, hence removing the need of using lexicons and decision trees which are the building blocks in [9], [60], [61]. This is one step toward removing expert knowledge when building an ASR system. Another advantage of character-based CTC is that it is more robust to the accented speech as the graphoneme sequence of words is less affected by accents than the phoneme

pronunciation [66]. Other output units that are larger than characters but smaller than words have also been studied [67].

It is a design challenge to determine the basic output unit to use for CTC prediction. In all the aforementioned works, the decomposition of a target word sequence into a sequence of basic units is fixed. However, the pre-determined fixed decomposition is not necessarily optimal. In [68], gram-CTC was proposed to automatically learn the most suitable decomposition of target sequences. Gram-CTC is based on characters, but allows to output variable number of characters (i.e., gram) at each time step. This not only boosts the modeling flexibility but also improves the final ASR system accuracy. However, all these works cannot be claimed as pure end-to-end systems because of the use of language models and decoders.

As the goal of ASR is to generate a word sequence from the speech waveform, word unit is the most natural output unit for network modeling. In [60], CTC with word output targets was explored but the accuracy is far from the phoneme-based CTC system. In [18], it was shown that by using 100k words as the output targets and by training the model with 125k hours of data, the CTC system with word units can beat the CTC system with phoneme unit. Fig. 1 gives an example of the posterior output of word CTC. In the figure, the units with the maximum posterior values are blanks and silences at most of time steps. All other posterior spikes come from word units. Hence, the ASR task becomes very simple: the output word sequence is constructed by taking the words corresponding to posterior spikes. No language model or complex decoding process is involved. Therefore, this can be considered as the first pure end-to-end ASR system whose success was built on top of very large amount of training data.




Fig. 1. An example of word CTC.

Compared to traditional cross-entropy training of LSTM, CTC is harder to train. First, the network initialization is very important. In [9], the LSTM network for CTC training was initialized from the LSTM network trained with cross entropy criterion. This can be circumvented by using very large amount of training data which also helps to prevent overfitting [60]. If the CTC network is randomly initialized, when presented very difficult samples, the CTC network tends to be very hard to train. In [15], a learning strategy called SortaGrad was proposed by first presenting the CTC network with shorter utterances (easy samples) and then longer utterances (hard samples) in the first training epoch. In the later epochs,

the utterances are given to CTC network randomly. This significantly improves the convergence of CTC training.

The spike patterns in Fig. 1 is general to CTC models with any modeling units. At the time steps where blank symbol dominates, it may be redundant to do the search as no information is provided. Given this observation, phone synchronous decoding [69] was proposed by skipping the search of blank-dominated time steps during CTC decoding. 2–3 times speedup was obtained without accuracy loss.

Note that the occurrence of spikes in CTC usually has a delay compared to the ground-truth location of the symbol. Such a delay introduces latency during the runtime decoding which is not desirable to the systems with realtime requirement. Therefore, a delay-constrained training was proposed in [61] by restricting the search paths used in the forward-backward process during CTC training to those in which the delay between CTC labels and the ground-truth alignment does not exceed a threshold. This constraint degrades the CTC performance a little, but the loss was recovered after sequence discriminative training.

Inspired by the CTC work, lattice-free maximum mutual information (LFMMI) [70] was recently proposed to train deep networks from scratch without initializing from cross-entropy networks. This single-step training has great advantage over current popular two-step training: first cross-entropy training and then sequence training. Lots of techniques have been developed to make LFMMI work, including a topology that the first frame of a phoneme has a different label than the remaining frames; a phoneme n-gram language model used to create denominator graph; a time-constraint similar to the delay-constrain used in CTC; several regularization methods to reduce overfitting; stacking multiple input frames as what CTC does; etc. LFMMI has been proven effective on tasks with different scales and underlying models.

Overall, there is clearly a major AM developing path from DNN to LSTM (temporal modeling) and then to CTC (end-to-end modeling). Although some modeling techniques, such as LFMMI, can achieve similar performance as CTC when phoneme is used as the modeling unit, they may not fit the trend of end-to-end modeling very well as these models require expert knowledge to design and need components such as language model and lexicon to work.

C. Attention-based Sequence-to-sequence Translation Models

Attention-based sequence-to-sequence model is another end-to-end model [71], [72]. It roots from the successful model in machine learning [73], [74] which extends the encoder-decoder framework [75] with an attention decoder. The attention model calculates the probability of the sequence \mathbf{l} as

$$P(\mathbf{l}|\mathbf{x}) = \prod_i P(l_i|\mathbf{x}\mathbf{l}_{1:i-1}) \quad (6)$$

with the probability at step i as

$$P(l_i|\mathbf{x}, \mathbf{l}_{1:i-1}) = \text{AttentionDecoder}(\mathbf{h}, \mathbf{l}_{1:i-1}) \quad (7)$$

$$\mathbf{h} = \text{Encoder}(\mathbf{x}). \quad (8)$$

The training criterion is to minimize $-\ln P(\mathbf{l}^*|\mathbf{x})$, where \mathbf{l}^* is the ground-truth labels sequence.

Different from the encoder in [75] which only takes the hidden vector of the last time step, the encoder in (8) transforms the whole speech input sequence \mathbf{x} to a high-level hidden vector sequence $\mathbf{h} = (\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_L), L \leq T$. Then, at each step in generating an output label l_i , an attention mechanism in (7) selects/weights the hidden vector sequence \mathbf{h} so that the most related hidden vectors are used for the prediction. Comparing (6) with (4) and (5), we can see the attention-based model does not have the frame-independence assumption imposed by CTC, which is the advantage of the attention model.

The attention-based model is even harder to train than the CTC model. There are plenty of tricks to be applied. For example, the vanilla attention-based model is highly complex during training if all the hidden vectors at all time steps are used in (7). Therefore, windowing method was used in [71] to reduce the number of candidates used in attention decoder. In [72], a pyramid structure was used in the encoder network so that only L high-level hidden vectors are generated instead of T hidden vectors from all the input time steps. Due to the high complexity and slow speed of training, the majority of attention-based works were only done at Univ. of Montreal [71] and Google [72], compared to the CTC works reported from many sites. However, because attention-model does not have the frame-independence assumption and may learn the implicit language model better, it was reported with better accuracy than CTC especially in the end-to-end setup without using external language models [76], [77].

The frame-independence assumption in CTC is the most criticized assumption as speech frames are correlated. On the other hand, the attention-based model has its drawback of not having monotonic left-to-right alignment and slow convergence. In [76], the attention training is combined with CTC training in a multi-task learning way by using the CTC objective function as the auxiliary cost function. Such a training strategy greatly improves the convergence of attention-based model and mitigates the alignment issue.

Although there are still arguments which end-to-end method is better for acoustic modeling, it seems at least now CTC clearly wins attention-based modeling since the end-to-end CTC model in [18] trained with hundreds of thousands of hours data has outperformed the traditional CTC models with context-dependent phone as target while the attention-based model still struggles to beat the traditional hybrid model even after very large improvement proposed recently [77] due to the difficulty of training.

IV. ACOUSTIC MODEL ROBUSTNESS

Current state of the art systems can achieve remarkable recognition accuracy when the test and training sets match, esp. when both under quiet close-talk condition. However, the performance dramatically degrades under mismatched or complicated environments such as higher noisy condition, including music or interfering talkers, or speech with strong accents [78], [79]. The solutions to this problem include adaptation, speech enhancement, and robust modeling.

A. Acoustic Model Adaptation

In this section, we use speaker adaptation as an example scenario to describe acoustic model adaptation technologies. The same technology should be easily applied to the adaptation of new environments and tasks, etc. Typically the speaker independent (SI) models are trained from a large dataset with objective to work best for all speakers. Speaker adaptation can significantly boost the performance of an individual speaker [80], [81]. However, we typically have limited adaptation data, and unsupervised adaptation is the main stream given prohibitive transcription cost. Current research focus is unsupervised adaptation with limited amount of speaker-dependent data, which can be addressed with better adaptation criterion and model topology. Since the adapted models are speaker dependent (SD), the size of the SD parameters is critical if we want to scale to millions of speakers. This requires solutions to minimizing the SD model footprint while maintaining the adaptation benefits.

Given the limited amount of adaptation data, the SD model should not be far away from the SI model. Reference [82] adds Kullback-Leibler divergence (KLD) regularization to the training criterion to prevent the adapted model from straying too far away from the SI model. This KLD adaptation criterion has been proven very effective in dealing with limited adaptation data. Most state-of-the-art SI models use senone (tied triphone states) as the output units. When limited amount of adaptation data is available, only very small amount of senones have been observed. In such a case, the adaptation turns to overfit the data distribution of these senones thus cannot generalize very well. In [83], a multi-task learning (MTL) framework was proposed by adding auxiliary monophone classification as the second task in addition to the primary senone classification task. As a result, the network adaptation is backed off to improving monophone classification accuracy when senones are not observed, hence increasing the generalization ability.

In contrast to adjusting the adaptation criterion, most of works focus on how to use very small amount of parameters to represent speaker characteristics. One solution is the singular value decomposition (SVD) bottleneck adaptation [84] which produces low-footprint SD models by making use of the SVD-restructured topology [85]. The linear transformation is applied to each of the bottleneck layer by adding a kXk SD matrices. The advantage of this approach is that only a couple of small matrices need to be updated for each speaker as k is the low-rank value of the SVD reconstruction and usually is very small. This dramatically reduces the deployment cost for speaker personalization while producing more reliable estimate of the adapted model [84]. Works have been done to further reduce the size of the kXk SD matrices. For example, when the adaptation data is very limited, the kXk matrix can be reduced to a diagonal matrix, such as learning hidden unit contribution (LHUC) [86], [87] and sigmoid adaptation [88]. This is a tradeoff between the modeling capacity and generalization. The LHUC and sigmoid adaptation have much smaller number of adaptation parameters compared to the SVD adaptation, but they may not get similar accuracy improvement when the amount of adaptation data is increased. The observation that

$k \times k$ SD matrices usually are diagonally dominant matrices inspired the proposal of low-rank plus diagonal (LRPD) decomposition which decomposes the $k \times k$ SD matrices into a diagonal matrix plus the multiplication of two low-rank matrices. By varying the low-rank values, the LRPD matrix generalizes the full-rank and the diagonal adaptation matrix, and hence can automatically utilize the adaptation data well instead of making tradeoff between model capacity and generalization.

The subspace methods are another type of methods that also aim to find a low dimensional subspace of the transformations, so that each transformation can be specified by a small number of parameters. One popular method in this category is the use of auxiliary features, such as i-vector [89], [90], speaker code [91], and noise estimate [92] which are concatenated with the standard acoustic features. It can be shown that the augmentation of auxiliary features is equivalent to confining the adapted bias vectors into a speaker subspace [93]. Furthermore, networks can be used to transform speaker features such as i-vectors into a bias to offset the speech feature into a speaker-normalized space [94]. In addition to augmenting features in the input space, the acoustic-factor features can also be appended in any layer of deep networks [95].

Other subspace methods include cluster adaptive training (CAT) [96], [97] and factorized hidden layer (FHL) [98], [99], where the transformations are confined into the speaker subspace. Similar to the eigenvoice [100] or cluster adaptive training [101] in the Gaussian mixture model era, CAT [96], [97] in DNN training constructs multiple DNNs to form the bases of a canonical parametric space. During adaptation, an interpolation vector which is associated to a target speaker or environment is estimated online to combine the multiple DNN bases into a single adapted DNN. Because only the combination vector is estimated, the adaptation only needs very small amount of data for fast adaptation. However, this is again a tradeoff from the model capacity. In contrast with online estimation of the combination vector, [102], [103] directly uses the posterior vectors of the acoustic context to enable fast unsupervised adaptation. The acoustic context factor can be speaker, gender, or acoustic environments such as noise and reverberation. The posterior calculation can be either independent [102] or dependent [103] on the recognizer network.

An issue in the CAT-style methods is that the bases are full-rank matrices, which require very large amount of training data. Therefore, the number of bases in CAT is usually constrained to a few [96], [97]. A solution is to use FHL [98], [99] which constrains the bases to be rank-1 matrices. In such a way, the training data for each basis is significantly reduced, enabling the use of larger number of bases. Also, FHL initializes the combination vector from i-vector for speaker adaptation, which helps to give the adaptation a very good starting point. In [104], LRPD was extended into the subspace-based approach to further reduce the speaker-specific footprint in a very similar way to FHL.

B. Speech Enhancement and Separation

It is well known that the current ASR systems perform poorly when the speech is corrupted with heavy noise or interfering speech [105], [106]. Although human listeners also suffer from poor audio signals, the performance degradation is significantly smaller than that in ASR systems.

In recent years, many works have been done to enhance speech under these conditions. Although majority of the works are focused on the single-channel speech enhancement and separation, the same techniques can be easily extended to multi-channel signals.

In the monaural speech enhancement and separation tasks, it is assumed that a linearly mixed single-microphone signal $y[n] = \sum_{s=1}^S x_s[n]$ is known and the goal is to recover the S streams of audio sources $x_s[n]$, $s = 1, \dots, S$. If there are only two audio sources, one for speech and one for noise (or music, etc.) and the goal is to recover the speech source, it's often called speech enhancement. If there are multiple speech sources, it is often referred to as speech separation. The enhancement and separation is usually carried out in the time-frequency domain, in which the task can be cast as recovering the short-time Fourier transformation (STFT) of the source signals $X_s(t, f)$ for each time frame t and frequency bin f , given the STFT of the mixed speech $Y(t, f) = \sum_{s=1}^S X_s(t, f)$.

Obviously, given only the mixed spectrum $Y(t, f)$, the problem of recovering $X_s(t, f)$ is under-determined (or ill-posed), as there are an infinite number of possible $X_s(t, f)$ combinations that lead to the same $Y(t, f)$. To overcome this problem, the system has to learn a model based on some training set \mathbb{S} that contains parallel sets of mixtures $Y(t, f)$ and their constituent target sources $X_s(t, f)$, $s = 1, \dots, S$ [20], [21], [107]–[112].

Over the decades, many attempts have been made to attack this problem. Before the deep learning era, the most popular techniques include computational auditory scene analysis (CASA) [113]–[115], non-negative matrix factorization (NMF) [116]–[118], and model based approach [119]–[121], such as factorial GMM-HMM [122]. Unfortunately these techniques only led to very limited success.

Recently, researchers have developed many deep learning techniques for speech enhancement and separation. The core of these techniques is to cast the enhancement or separation problem into a supervised learning problem. More specifically, the deep learning models are optimized to predict the source belonging to the target class, usually for each time-frequency bin, given the pairs of (usually artificially) mixed speech and source streams. Compared to the original setup of unsupervised learning, this is a significant step forward and leads to great progress in speech enhancement. This simple strategy, however, is still not satisfactory, as it only works for separating audios with very different characteristics, such as separating speech from (often challenging) background noise (or music) or speech of a specific speaker from other speakers [110]. It does not work well for speaker-independent multi-talker speech separation.

The difficulty in speaker-independent multi-talker speech separation comes from the label ambiguity or permutation

problem. Because audio sources are symmetric given the mixture (i.e., $x_1 + x_2$ equals to $x_2 + x_1$ and both x_1 and x_2 have the same characteristics), there is no pre-determined way to assign the correct source target to the corresponding output layer during supervised training. As a result, the model cannot be well trained to separate speech.

Fortunately, several techniques have been proposed to address the label ambiguity problem [20], [21], [106], [111], [112], [123]. In Weng *et al.* [106] the instantaneous energy was used to solve the label ambiguity problem and a two-speaker joint-decoder with a speaker switching penalty was used to separate and trace speakers. This work achieved the best result on the dataset used in 2006 monaural speech separation and recognition challenge [105]. However, energy, which is manually picked, may not be the best information to assign labels in all conditions. Actually, we have found that in many cases pitch difference is a more important cue. In Hershey *et al.* [111], [112] a novel technique called deep clustering (DPCL) was proposed. In this model, it is assumed that each time-frequency bin belongs to only one speaker. During training, each time-frequency bin is mapped into an embedding space. The embedding is optimized so that time-frequency bins belong to the same speaker are closer and that of different speakers are farther away in this space. During evaluation, a clustering algorithm is used upon embeddings to generate a partition of the time-frequency bins. To further improve the performance, they stacked yet another network to estimate real masks for each source stream given the results from the deep clustering [112]. Chen *et al.* [123] proposed a technique called deep attractor network (DANet). Following DPCL, their approach also learns a high-dimensional embedding of the acoustic signals. Different from DPCL, however, it creates cluster centers, called attractor points, in the embedding space to pull together the time-frequency bins corresponding to the same source. The main limitation of DANet is the requirement to estimate attractor points during evaluation time.

In Yu *et al.* [20] and Kolbæk *et al.* [21], a simpler technique named permutation invariant training (PIT) was proposed to attack the speaker independent multi-talker speech separation problem. In this new approach, the source targets are treated as a set (i.e., order is irrelevant). During training, PIT first determines the output-target assignment with the minimum error at the utterance level based on the forward-pass result. It then minimizes the error given the assignment. This strategy elegantly solves the label permutation problem and speaker tracing problem in one shot. Unlike other techniques such as DPCL and DANet that require a separate clustering step to trace speech streams during evaluation, PIT does not require a separate tracing step (and thus can be used in real-time systems). Instead, each output layer is corresponding to one stream of sources. In PIT the computational cost associated with label assignment is negligible compared to the network forward computation during training, and no label assignment (and thus no cost) is needed during evaluation. Recently, Hershey *et al.*¹ have found out that in DPCL the embeddings actually are grouped into two classes, instead of many different

classes for different speakers, in the two-speaker separation problem. This indicates that DPCL essentially learns separation models that is very similar to that learned by PIT. DPCL, DANet, and PIT all achieve similar performance on speaker-independent two- to three-talker speech separation tasks yet PIT is the simplest among all, can be used in real-time systems, and can be easily combined with other techniques. Moreover, unlike DPCL or DaNet, PIT does not need to know or estimate the number of streams in the mixture. We therefore believe that PIT is most promising among these techniques. Similar to progresses made when converting the speech separation problem from an unsupervised learning problem into a supervised learning problem, PIT converts the speech separation problem from supervised learning with ambiguous labels to that with clear labels.

For speech recognition, we can feed each separated speech stream to ASR systems. Even better, the deep learning based AM may be jointly optimized end-to-end with the separation component, which is often an RNN. Since separation is just an intermediate step, Yu *et al.* [124] proposed to directly optimize the cross-entropy criterion against senone labels using PIT without having an explicit speech separation step. Their preliminary results on AMI dataset indicate that PIT can significantly improve the recognition accuracy, compared to the models trained with single-talker speech, when recognizing two-talker mixed speech.

C. Robust Training

The success of deep neural networks relies on the availability of a large amount of transcribed data to train millions of model parameters. However, deep models still suffer reduced performance when exposed to test data from a new domain. Because it is typically very time-consuming or expensive to transcribe large amounts of data for a new domain, domain-adaptation approaches have been proposed to bootstrap the training of a new system from an existing well-trained model [81], [84]. These methods still require transcribed data from the new domain and thus their effectiveness is limited by the amount of transcribed data available in the new domain. Although unsupervised adaptation methods can be used by generating labels from a decoder, the performance gap between supervised and unsupervised adaptation is large [81].

Recently, the concept of adversarial training [125] was explored for noise-robust ASR [126]–[128]. This solution is a pure unsupervised domain adaptation method without utilizing too much knowledge about the new domain. The idea is to have three networks in the model: the encoder network, the recognizer network, and the domain discriminator network. The encoder network generates the intermediate representation, which will be used in both the recognizer network to generate posteriors of phoneme units and the domain discriminator network to generate domain labels. The intermediate representation is learned adversarially to the domain discriminator, i.e., to minimize the domain classification accuracy. In such a way, the intermediate representation is invariant to the input of different domains. At the same time, the intermediate representation is trained to maximize

¹based on personal communication

the phoneme classification accuracy with the source domain labels. During testing time, the encoder network generates the intermediate representation from the target domain data and input it into the recognizer network. The training is done by inserting a gradient reverse layer (GRL) [129] between the encoder network and the domain discriminator network. During forward propagation, GRL acts as an identity transform. During back propagation, GRL takes the gradient from the domain discriminator network, multiplies it by a negative constant, and then passes it to the encoder network.

The advantage of GRL-based unsupervised adaptation is that it does not require any knowledge about the target domain. In contrast, the adaptation to the new domain should be more effective if we have some domain knowledge by simulating the target domain data. For example, if the source domain is a clean environment and the target domain is a noisy environment, we can simulate noisy environment data and then do the multi-style training [130] with simulated data, e.g., in [131]. However, the multi-style training does not use the well-trained source model which usually has very high accuracy in the source domain.

Recently, the teacher/student (T/S) learning [132] method was proposed to perform adaptation without the use of transcriptions. The data from the source domain are processed by the source-domain model (teacher) to generate the corresponding posterior probabilities or soft labels. These posterior probabilities are used in lieu of the usual hard labels derived from the transcriptions to train the target (student) model with the parallel data from the target domain. With this approach, the network can be trained on a potentially enormous amount of training data and the challenge of adapting a large-scale system shifts from transcribing thousands of hours of audio to the potentially much simpler and lower-cost task of designing a scheme to generate the appropriate parallel data. Evaluated on the CHiME-3 task, the T/S learning method can get 40+ % relative WER reduction with only several thousands paired utterances [132], much larger than what can be obtained from the traditional feature mapping and mask learning methods.

The T/S learning approach is closely related to other approaches for adaptation or retraining that employ knowledge distillation [133]. In these approaches, the soft labels generated by a teacher model are used as a regularization term to train a student model with conventional hard labels. For example, knowledge distillation was used to train a system on the Aurora 2 noisy digit recognition task, using the clean and noisy training sets [134]. In [135] it was shown that for the multi-channel CHiME-4 task, soft labels could be derived using enhanced features generated by a beamformer then processed through a network trained with conventional multi-style training. In all cases, the soft labels provided by the teacher network regularized the conventional training of the student network using hard labels derived from transcriptions. Thus, the use of additional unlabeled training data was not possible.

In conclusion, domain adaption without labeled data will be an important research direction. If we do not have any knowledge about the target domain, adversarial training should be a good way to go. On the other hand, if we can simulate

data similar to the target domain data, T/S learning and knowledge distillation are good methods. Especially, the T/S learning methods forgoes the need for hard labels from the data in the new domain entirely and relies solely on the soft labels provided by the parallel corpus and well-trained source model. This allows the use of a significantly larger set of adaptation data which adds robustness to the resulting model.

V. ACOUSTIC MODELS WITH EFFICIENT DECODING

Training deep networks by stacking multiple layers helps to improve WER. However, the computational cost becomes a concern, especially to the industry deployment where realtime is always with high priority. There are several ways to reducing the runtime cost.

The first one is to use singular value decomposition (SVD) which was originally proposed in [85] and has been widely used. The SVD method decomposes a full-rank matrix into two lower-rank matrices, hence can significantly reduce the number of parameters in deep models without losing accuracy after retraining. This is general to any deep network structure. In [136], [137], a similar method was proposed for learning compact LSTMs via low-rank factorization and parameter sharing schemes.

The second way is to employ teacher-student (T/S) learning or knowledge distillation, such as the works in [138]–[141]. T/S learning was proposed in [142] to compress a standard DNN model by minimizing the KLD between the output distributions of the small-size DNN and a standard large-size DNN. The learning equals to the CE training using the soft label generated by the teacher model as the target for the student learning. The concept of T/S learning was extended as the concept of distilling the knowledge in [133] by combining the CE training using soft labels with the standard CE training using the 1-hot vector as the target. The soft target in knowledge distillation serves as the regularization term to the standard CE training. Recently, Microsoft and IBM continuously broke the WER record on the Switchboard task [143], [144]. The built systems are usually giant models ensemble of multiple deep models. Such system cannot be deployed to realtime application. In such a scenario, T/S learning or knowledge distillation provides a good solution to getting a compact model with high modeling capacity.

The third method is to compress the models by heavy quantization, applying either very low-bit quantization or vector quantization. Reference [145] gives very nice summary of technologies to speed up the runtime evaluation of deep networks. Those technologies including 8-bit quantization do not require re-training deep networks. However, the ASR accuracy is significantly reduced when the model is compressed heavily into even lower bits or the network structure becomes more complex. Therefore, refining with quantization is important for both very low-bit quantization [146], [147] and vector quantization [148] so that the training and testing are consistent.

The fourth solution is to work on model structures. LSTM with projection layer (LSTMP) was proposed to reduce the computational cost by adding a linear projection layer after the

LSTM layer [8]. This projected vector has lower dimension than the output of the LSTM layer, and is used to replace the recurrent input h_{t-1} to the LSTM unit. Although the initial purpose of LSTMP is to reduce runtime cost, it was also reported helpful to error rate reduction [8], [26] which is possibly because parameter reduction helps to generalize better if the LSTM model is too strong. The projection layer is not used in the later Google's CTC work [9], [60], [61] where the number of memory cells in LSTM is smaller and the amount of training data is 20 times larger. In [27], the runtime cost of LSTM is reduced with model simplification by coupling input gates and forget gates.

Finally, the correlation across frames can be used to reduce the evaluation frequency of deep network scores. For DNNs or CNNs, this can be done with a frame-skipping strategy by computing the acoustic scores once every several frames and copying the acoustic scores to the frames in which no acoustic score is evaluated during decoding [149]. However, for decoding with LSTM models, the same frame-skipping strategy needs to be done in both training and testing stages so that the behaviors of LSTM's memory units are consistent [27]. Recently, LFMMI [70] and lower frame rate (LFR) LSTM-RNN [150] were proposed to work on the 30 ms units instead of the traditional 10 ms units by modeling phonemes instead of states. Because of using larger input frame skips, the LFMMI and LFR models only need 1/3 computation cost of the traditional networks which operate on 10 ms inputs. As CTC uses phones or even words as the output targets, it can also work on 30 ms or even larger inputs, therefore significantly reduce the runtime cost.

VI. FUTURE DIRECTIONS

The research frontier has been shifted from ASR with close-talk microphones to ASR with far-field microphones, driven by increased demand from users to interact with devices without wearing or carrying a close-talk microphone. For example, Amazon's Echo and Google Home have now been deployed in many families across the world. Many difficulties hidden in close-talk scenarios now surface when far-field microphones are used. This is because in the far-field scenario, the energy of speech signal is very low when it reaches the microphones. Comparatively, the interfering signals, such as background noise, reverberation, and speech from other talkers, become so distinct that they can no longer be ignored.

Although many of the speech recognition techniques developed for close-talk scenarios can be directly applied to far-field scenarios, these techniques show inferior performance under the distant recognition scenario. To ultimately solve the distant speech recognition problem, we need to optimize the whole pipeline starting from audio capturing (e.g., microphone-array signal processing) to acoustic modeling and decoding. We perceive following possible research directions.

First, although we have made some interesting progresses in monaural speech enhancement, separation, and recognition, much more improvements are desired. For example, the current PIT system performs very well when separating speech mixtures of different genders. However, the separation quality

is deteriorated when two same-gender speakers speak at the same time. Will the quality be good enough if we also exploit beam-forming results and multi-channel information? Is there a better model than conventional LSTMs for speech separation and tracing? How can we exploit additional information, such as language model, by feeding information from the decoder back to the speech enhancement and separation component and by jointly considering all streams of speech when making decoding decision?

Second, the end-to-end optimization strategy is desired, given its simplicity and joint optimization characteristics, if we only need to optimize for the decoding result and have sufficient training data. This has been proven effective with word-based CTC when trained with hundreds of thousands hours of data. However, it is not feasible to get that large amount of data for most tasks. Considering that current end-to-end systems trained with thousands hours of data use two LMs, one that is implicitly built-in and trained with the AM, and one that is separated and trained using text data, we would guess that further accuracy improvement can be achieved if we can integrate the second LM into the end-to-end system and optimize them jointly when audio signals are available and separately when only text data is available. If this turns out to be beneficial, would it affect our choice of modeling unit in end-to-end systems? For example, if Chinese characters are used as modeling unit, adding a new character would be difficult, partly because it will change the number of classes in the model and partly because there may not be enough data to train the new character. However, since the pronunciation of the new character would be one of the syllables that have already been covered by other characters, adding new characters can be extremely easy if syllable is the modeling unit since the pronunciation knowledge can be directly transferred. To further extend it, can we design the model so that components are built and can be transferred at different scales? Both the CTC and attention-based model have their individual pros and cons. The joint training with two models together is first but superficial step toward better end-to-end modeling. Can we formulate a single model by taking the advantages of both models?

Third, even the models trained with huge amount of data are lack of robustness. This is because training-test mismatch is unavoidable given the cost of data collection. Adding simulated data can alleviate the problem if we can foresee the possible variations but may still be not sufficient. Can we design a model (e.g., as a special nonlinear dynamic system) that constantly adapts itself within some limit, e.g., controlled by some kernel size? Can such models automatically exploit information gained from past similar speakers to quickly adapt to new speakers? It is even more interesting if such models can gradually identify the most reliable regularities in the data.

REFERENCES

- [1] D. Yu, L. Deng, and G. E. Dahl, "Roles of pre-training and fine-tuning in context-dependent DBN-HMMs for real-world speech recognition," in *Proc. NIPS 2010 Workshop on Deep Learning and Unsupervised Feature Learning*, 2010.

- [2] G. E. Dahl, D. Yu, L. Deng, and A. Acero, "Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition," *IEEE Trans. Audio Speech Lang. Processing*, vol. 20, no. 1, pp. 30–42, Jan. 2012.
- [3] D. Yu, F. Seide, and G. Li, "Conversational speech transcription using context-dependent deep neural networks," in *Proc. 29th Int. Conf. Int. Conf. Machine Learning*, Edinburgh, Scotland, 2011, pp. 437–440.
- [4] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A. R. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, and B. Kingsbury, "Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups," *IEEE Signal Processing Mag.*, vol. 29, no. 6, pp. 82–97, Nov. 2012.
- [5] O. Abdel-Hamid, A. R. Mohamed, H. Jiang, and G. Penn, "Applying convolutional neural networks concepts to hybrid NN-HMM model for speech recognition," in *Proc. 2012 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Kyoto, Japan, 2012, pp. 4277–4280.
- [6] L. Deng, J. Li, J. T. Huang, K. S. Yao, D. Yu, F. Seide, M. Seltzer, G. Zweig, X. D. He, J. Williams, Y. F. Gong, and A. Acero, "Recent advances in deep learning for speech research at microsoft," in *Proc. 2013 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Vancouver, BC, Canada, 2013, pp. 8604–8608.
- [7] 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 Lang. Processing*, vol. 22, no. 10, pp. 1533–1545, Oct. 2014.
- [8] H. Sak, A. Senior, and F. Beaufays, "Long short-term memory recurrent neural network architectures for large scale acoustic modeling," in *15th Proc. Interspeech*, Singapore, Singapore, 2014, pp. 338–342.
- [9] H. Sak, A. Senior, K. Rao, O. Irsoy, A. Graves, F. Beaufays, and J. Schalkwyk, "Learning acoustic frame labeling for speech recognition with recurrent neural networks," in *Proc. 2015 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Brisbane, QLD, Australia, 2015, pp. 4280–4284.
- [10] T. N. Sainath, O. Vinyals, A. Senior, and H. Sak, "Convolutional, long short-term memory, fully connected deep neural networks," in *Proc. 2015 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Brisbane, QLD, Australia, 2015, pp. 4580–4584.
- [11] M. X. Bi, Y. M. Qian, and K. Yu, "Very deep convolutional neural networks for LVCSR," in *16th Proc. Interspeech*, Dresden, Germany, 2015, pp. 3259–3263.
- [12] V. Mitra and H. Franco, "Time-frequency convolutional networks for robust speech recognition," in *Proc. 2015 IEEE Workshop on Automatic Speech Recognition and Understanding*, Scottsdale, AZ, USA, 2015, pp. 317–323.
- [13] V. Peddinti, D. Povey, and S. Khudanpur, "A time delay neural network architecture for efficient modeling of long temporal contexts," in *16th Proc. Interspeech*, Dresden, Germany, 2015, pp. 3214–3218.
- [14] T. Sercu, C. Puhresch, B. Kingsbury, and Y. LeCun, "Very deep multilingual convolutional neural networks for LVCSR," in *Proc. 2016 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Shanghai, China, 2016, pp. 4955–4959.
- [15] D. Amodei, R. Anubhai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, J. D. Chen, M. Chrzanowski, A. Coates, G. Diamos, E. Elsen, J. Engel, L. X. Fan, C. Fougner, T. Han, A. Hannun, B. Jun, P. LeGresley, L. Lin, S. Narang, A. Ng, S. Ozair, R. Prenger, J. Raiman, S. Satheesh, D. Seetapun, S. Sengupta, Y. Wang, Z. Q. Wang, C. Wang, B. Xiao, D. N. Yogatama, J. Zhan, and Z. Y. Zhu, "Deep speech 2: End-to-end speech recognition in English and Mandarin," arXiv:1512.02595, 2015.
- [16] S. L. Zhang, C. Liu, H. Jiang, S. Wei, L. R. Dai, and Y. Hu, "Feedforward sequential memory networks: A new structure to learn long-term dependency," arXiv:1512.08301, 2015.
- [17] D. Yu, W. Xiong, J. Droppo, A. Stolcke, G. L. Ye, J. Li, and G. Zweig, "Deep convolutional neural networks with layer-wise context expansion and attention," in *17th Proc. Interspeech*, San Francisco, USA, 2016.
- [18] H. Soltau, H. Liao, and H. Sak, "Neural speech recognizer: acoustic-to-word LSTM model for large vocabulary speech recognition," arXiv:1610.09975, 2016.
- [19] W. Xiong, J. Droppo, X. Huang, F. Seide, M. Seltzer, A. Stolcke, D. Yu, and G. Zweig, "Achieving human parity in conversational speech recognition," arXiv:1610.05256, 2016.
- [20] D. Yu, M. Kolbæk, Z. H. Tan, and J. Jensen, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," arXiv:1607.00325, 2017.
- [21] M. Kolbæk, D. Yu, Z. H. Tan, and J. Jensen, "Multi-talker speech separation and tracing with permutation invariant training of deep recurrent neural networks," arXiv:1703.06284, 2017.
- [22] D. Yu and L. Deng, *Automatic Speech Recognition: A Deep Learning approach*. London: Springer, 2015.
- [23] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
- [24] A. Graves, A. R. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," in *Proc. 2013 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Vancouver, BC, Canada, 2013, pp. 6645–6649.
- [25] X. G. Li and X. H. Wu, "Constructing long short-term memory based deep recurrent neural networks for large vocabulary speech recognition," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Brisbane, QLD, Australia, 2015, pp. 4520–4524.
- [26] Y. J. Miao and F. Metze, "On speaker adaptation of long short-term memory recurrent neural networks," in *16th Proc. Interspeech*, Dresden, Germany, 2015, pp. 1101–1105.
- [27] Y. J. Miao, J. Li, Y. Q. Wang, S. X. Zhang, and Y. F. Gong, "Simplifying long short-term memory acoustic models for fast training and decoding," in *Proc. 2016 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Shanghai, China, 2016.
- [28] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, "Empirical evaluation of gated recurrent neural networks on sequence modeling," arXiv:1412.3555, 2014.
- [29] Y. Zhang, G. G. Chen, D. Yu, K. S. Yao, S. Khudanpur, and J. Glass, "Highway long short-term memory RNNs for distant speech recognition," in *Proc. 2016 IEEE International Conference on Acoustics, Speech and Signal Processing*, Shanghai, China, 2016.
- [30] Y. Y. Zhao, S. Xu, and B. Xu, "Multidimensional residual learning based on recurrent neural networks for acoustic modeling," in *17th Proc. Interspeech*, San Francisco, USA, 2016, pp. 3419–3423.
- [31] J. Kim, M. El-Khany, and J. Lee, "Residual LSTM: Design of a deep recurrent architecture for distant speech recognition," arXiv:1701.03360, 2017.
- [32] K. He, X. Y. Zhang, S. Q. Ren, and J. Sun, "Deep residual learning for image recognition," arXiv:1512.03385, 2015.
- [33] A. R. Mohamed, G. Hinton, and G. Penn, "Understanding how deep belief networks perform acoustic modelling," in *Proc. 2012 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Kyoto, Japan, 2012, pp. 4273–4276.
- [34] J. Li, D. Yu, J. T. Huang, and Y. F. Gong, "Improving wideband speech recognition using mixed-bandwidth training data in CD-DNN-HMM," in *Proc. 2012 IEEE Spoken Language Technology Workshop*, Miami, FL, USA, 2012, pp. 131–136.
- [35] J. Li, A. Mohamed, G. Zweig, and Y. F. Gong, "LSTM time and frequency recurrence for automatic speech recognition," in *Proc. 2015 IEEE Workshop on Automatic Speech Recognition and Understanding*, Scottsdale, AZ, USA, 2015.
- [36] J. Li, A. Mohamed, G. Zweig, and Y. F. Gong, "Exploring multidimensional LSTMs for large vocabulary ASR," in *Proc. 2016 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Shanghai, China, 2016.
- [37] T. N. Sainath and B. Li, "Modeling time-frequency patterns with LSTM vs. convolutional architectures for LVCSR tasks," in *17th Proc. Interspeech*, San Francisco, USA, 2016.
- [38] N. Kalchbrenner, I. Danihelka, and A. Graves, "Grid long short-term memory," arXiv:1507.01526, 2015.
- [39] W. N. Hsu, Y. Zhang, and J. Glass, "A prioritized grid long short-term memory RNN for speech recognition," in *Proc. 2016 IEEE Spoken Language Technology Workshop (SLT)*, San Diego, California, USA, 2016, pp. 467–473.
- [40] A. Graves and J. Schmidhuber, "Framewise phoneme classification with bidirectional LSTM and other neural network architectures," *Neural Netw.*, vol. 18, no. 5–6, pp. 602–610, Jul.–Aug. 2005.
- [41] S. F. Xue and Z. J. Yan, "Improving latency-controlled BLSTM acoustic models for online speech recognition," in *Proc. 2017 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, New Orleans, USA, 2017.
- [42] Y. LeCun and Y. Bengio, "Convolutional networks for images, speech, and time-series," in *The Handbook of Brain Theory and Neural Networks*, M. A. Arbib, Ed. Cambridge: MIT Press, 1995.
- [43] K. J. Lang, A. H. Waibel, and G. E. Hinton, "A time-delay neural network architecture for isolated word recognition," *Neural Netw.*, vol. 3, no. 1, pp. 23–43, Dec. 1990.

- [44] O. Abdel-Hamid, L. Deng, and D. Yu, "Exploring convolutional neural network structures and optimization techniques for speech recognition," in *14th Proc. Interspeech*, Lyon, France, 2013, pp. 3366–3370.
- [45] T. N. Sainath, A. R. Mohamed, B. Kingsbury, and B. Ramabhadran, "Deep convolutional neural networks for LVCSR," in *Proc. 2013 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Vancouver, BC, Canada, 2013, pp. 8614–8618.
- [46] T. Sercu and V. Goel, "Dense prediction on sequences with time-dilated convolutions for speech recognition," arXiv:1611.09288, 2016.
- [47] L. Tóth, "Modeling long temporal contexts in convolutional neural network-based phone recognition," in *Proc. 2015 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Brisbane, QLD, Australia, 2015, pp. 4575–4579.
- [48] T. Zhao, Y. X. Zhao, and X. Chen, "Time-frequency kernel-based CNN for speech recognition," in *16th Proc. Interspeech*, Dresden, Germany, 2015.
- [49] N. Jaitly and G. Hinton, "Learning a better representation of speech soundwaves using restricted boltzmann machines," in *Proc. 2011 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Prague, Czech Republic, 2011, pp. 5884–5887.
- [50] D. Palaz, R. Collobert, and M. Magimai-Doss, "Estimating phoneme class conditional probabilities from raw speech signal using convolutional neural networks," in *14th Proc. Interspeech*, Lyon, France, 2014.
- [51] Z. Tüske, P. Golik, R. Schlüter, and H. Ney, "Acoustic modeling with deep neural networks using raw time signal for LVCSR," in *15th Proc. Interspeech*, Singapore, Singapore, 2014, pp. 890–894.
- [52] T. N. Sainath, R. J. Weiss, A. W. Senior, K. W. Wilson, and O. Vinyals, "Learning the speech front-end with raw waveform CLDNNs," in *16th Proc. Interspeech*, Dresden, Germany, 2015, pp. 1–5.
- [53] H. Dinkel, N. X. Chen, Y. M. Qian, and K. Yu, "End-to-end spoofing detection with raw waveform CLDNNs," in *Proc. 2017 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, New Orleans, USA, 2017.
- [54] T. Yoshioka, N. Ito, M. Delcroix, A. Ogawa, K. Kinoshita, M. Fujimoto, C. Z. Yu, W. J. Fabian, M. Espi, T. Higuchi, S. Araki, and T. Nakatani, "The NTT chime-3 system: Advances in speech enhancement and recognition for mobile multi-microphone devices," in *Proc. 2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, Scottsdale, AZ, USA, 2015, pp. 436–443.
- [55] X. Xiao, S. Watanabe, H. Erdogan, L. Lu, J. Hershey, M. L. Seltzer, G. G. Chen, Y. Zhang, M. Mandel, and D. Yu, "Deep beamforming networks for multi-channel speech recognition," in *Proc. 2016 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Shanghai, China, 2016, pp. 5745–5749.
- [56] T. N. Sainath, R. J. Weiss, K. W. Wilson, A. Narayanan, M. Bacchiani et al., "Speaker location and microphone spacing invariant acoustic modeling from raw multichannel waveforms," in *Proc. 2015 IEEE Int. Conf. Automatic Speech Recognition and Understanding (ASRU)*, Scottsdale, AZ, USA, 2015, pp. 30–36.
- [57] T. N. Sainath, R. J. Weiss, K. W. Wilson, A. Narayanan, and M. Bacchiani, "Factored spatial and spectral multichannel raw waveform CLDNNs," in *Proc. 2016 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Shanghai, China, 2016, pp. 5075–5079.
- [58] T. N. Sainath, R. J. Weiss, K. W. Wilson, B. Li, A. Narayanan, E. Variiani, M. Bacchiani, I. Shafran, A. Senior, K. W. Chin, A. Misra, and C. Kim, "Multichannel signal processing with deep neural networks for automatic speech recognition," *IEEE/ACM Trans. Audio Speech Language Processing*, vol. 25, no. 5, pp. 965–979, May 2017.
- [59] E. Variiani, T. N. Sainath, I. Shafran, and M. Bacchiani, "Complex linear projection (CLP): A discriminative approach to joint feature extraction and acoustic modeling," in *17th Proc. Interspeech*, San Francisco, USA, 2016, pp. 808–812.
- [60] H. Sak, A. Senior, K. Rao, and F. Beaufays, "Fast and accurate recurrent neural network acoustic models for speech recognition," in *16th Proc. Interspeech*, Dresden, Germany, 2015.
- [61] A. Senior, H. Sak, F. de Chaumont Quiry, T. Sainath, and K. Rao, "Acoustic modelling with CD-CTC-SMBR LSTM RNNs," in *Proc. 2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, Scottsdale, AZ, USA, 2015, pp. 604–609.
- [62] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proc. the 23rd Int. Conf. Machine Learning*, Pittsburgh, Pennsylvania, USA, 2006, pp. 369–376.
- [63] A. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Satheesh, S. Sengupta, A. Coates, and A. Y. Ng, "Deep speech: Scaling up end-to-end speech recognition," arXiv:1412.5567, 2014.
- [64] Y. Miao, M. Gowayyed, and F. Metze, "EESSEN: End-to-end speech recognition using deep RNN models and WFST-based decoding," in *Proc. 2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, Scottsdale, AZ, USA, 2015, pp. 167–174.
- [65] Y. J. Miao, M. Gowayyed, X. Y. Na, T. Ko, F. Metze, and A. Waibel, "An empirical exploration of etc acoustic models," in *Proc. 2016 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Shanghai, China, 2016, pp. 2623–2627.
- [66] K. Rao and H. Sak, "Multi-accent speech recognition with hierarchical grapheme based models," in *Proc. 2017 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, New Orleans, USA, 2017.
- [67] G. Zweig, C. Z. Yu, J. Droppo, and A. Stolcke, "Advances in all-neural speech recognition," in *Proc. 2017 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, New Orleans, USA, 2017.
- [68] H. R. Liu, Z. Y. Zhu, X. G. Li, and S. Satheesh, "Gram-CTC: Automatic unit selection and target decomposition for sequence labelling," arXiv:1703.00096, 2017.
- [69] Z. H. Chen, Y. M. Zhuang, Y. M. Qian, and K. Yu, "Phone synchronous speech recognition with etc lattices," *IEEE/ACM Trans. Audio, Speech, and Language Processing*, vol. 25, no. 1, pp. 90–101, 2017.
- [70] D. Povey, V. Peddinti, D. Galvez, P. Ghahramani, V. Manohar, X. Y. Na, Y. M. Wang, and S. Khudanpur, "Purely sequence-trained neural networks for ASR based on lattice-free MMI," in *17th Proc. Interspeech*, San Francisco, USA, 2016.
- [71] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, "End-to-end attention-based large vocabulary speech recognition," in *Proc. 2016 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Shanghai, China, 2016, pp. 4945–4949.
- [72] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *Proc. 2016 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Shanghai, China, 2016, pp. 4960–4964.
- [73] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," arXiv:1409.0473, 2014.
- [74] V. Mnih, N. Heess, A. Graves, and K. Kavukcuoglu, "Recurrent models of visual attention," in *Advances in Neural Information Processing Systems 27: 28th Annual Conference on Neural Information Processing Systems*, Montreal, Canada, 2014, pp. 2204–2212.
- [75] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," arXiv:1406.1078, 2014.
- [76] S. Kim, T. Hori, and S. Watanabe, "Joint CTC-attention based end-to-end speech recognition using multi-task learning," in *Proc. 2017 IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, New Orleans, USA, 2017.
- [77] Y. Zhang, W. Chan, and N. Jaitly, "Very deep convolutional networks for end-to-end speech recognition," in *Proc. 2017 IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, New Orleans, USA, 2017.
- [78] J. Li, L. Deng, Y. F. Gong, and R. Haeb-Umbach, "An overview of noise-robust automatic speech recognition," *IEEE/ACM Trans. Audio, Speech, and Language Processing*, vol. 22, no. 4, pp. 745–777, Apr. 2014.
- [79] J. Li, L. Deng, R. Haeb-Umbach, and Y. F. Gong, *Robust Automatic Speech Recognition: A Bridge to Practical Applications*, Waltham: Academic Press, 2015.
- [80] F. Seide, G. Li, X. Chen, and D. Yu, "Feature engineering in context-dependent deep neural networks for conversational speech transcription," in *Proc. 2011 IEEE Workshop on Automatic Speech Recognition and Understanding*, Waikoloa, HI, USA, 2011, pp. 24–29.
- [81] H. Liao, "Speaker adaptation of context dependent deep neural networks," in *Proc. 2013 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Vancouver, BC, Canada, 2013, pp. 7947–7951.
- [82] D. Yu, K. S. Yao, H. Su, G. Li, and F. Seide, "KL-divergence regularized deep neural network adaptation for improved large vocabulary speech recognition," in *Proc. 2013 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Vancouver, BC, Canada, 2013, pp. 7893–7897.
- [83] Z. Huang, J. Li, S. M. Siniscalchi, I. F. Chen, J. Wu, and C. H. Lee, "Rapid adaptation for deep neural networks through multi-task learning," in *16th Proc. Interspeech*, Dresden, Germany, 2015, pp. 3625–3629.
- [84] J. Xue, J. Li, D. Yu, M. Seltzer, and Y. F. Gong, "Singular value decomposition based low-footprint speaker adaptation and personalization

- deep neural network,” in *Proc. 2014 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Florence, Italy, 2014, pp. 6359–6363.
- [85] J. Xue, J. Li, and Y. F. Gong, “Restructuring of deep neural network acoustic models with singular value decomposition,” in *14th Proc. Interspeech*, Lyon, France, 2013, pp. 2365–2369.
- [86] P. Swietojanski and S. Renals, “Learning hidden unit contributions for unsupervised speaker adaptation of neural network acoustic models,” in *Proc. 2014 IEEE Spoken Language Technology Workshop*, South Lake Tahoe, NV, USA, 2014.
- [87] P. Swietojanski, J. Li, and S. Renals, “Learning hidden unit contributions for unsupervised acoustic model adaptation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 8, pp. 1450–1463, Aug. 2016.
- [88] Y. Zhao, J. Li, J. Xue, and Y. F. Gong, “Investigating online low-footprint speaker adaptation using generalized linear regression and click-through data,” in *Proc. 2015 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Brisbane, QLD, Australia, 2015, pp. 4310–4314.
- [89] G. Saon, H. Soltau, D. Nahamoo, and M. Picheny, “Speaker adaptation of neural network acoustic models using i-vectors,” in *Proc. 2013 IEEE Workshop on Automatic Speech Recognition and Understanding*, Olomouc, Czech Republic, 2013, pp. 55–59.
- [90] A. Senior and I. Lopez-Moreno, “Improving DNN speaker independence with i-vector inputs,” in *Proc. 2014 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Florence, Italy, 2014, pp. 225–229.
- [91] O. Abdel-Hamid and H. Jiang, “Fast speaker adaptation of hybrid NN/HMM model for speech recognition based on discriminative learning of speaker code,” in *Proc. 2013 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Vancouver, BC, Canada, 2013, pp. 7942–7946.
- [92] M. L. Seltzer, D. Yu, and Y. Q. Wang, “An investigation of deep neural networks for noise robust speech recognition,” in *Proc. 2013 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Vancouver, BC, Canada, 2013, pp. 7398–7402.
- [93] D. Yu and L. Deng, “Adaptation of deep neural networks,” in *Automatic Speech Recognition*, D. Yu and L. Deng, Eds. London: Springer, 2015, pp. 193–215.
- [94] Y. J. Miao, H. Zhang, and F. Metzger, “Towards speaker adaptive training of deep neural network acoustic models,” in *15th Proc. Interspeech*, Singapore, Singapore, 2014.
- [95] J. Li, J. T. Huang, and Y. F. Gong, “Factorized adaptation for deep neural network,” in *Proc. 2014 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Florence, Italy, 2014.
- [96] T. Tan, Y. M. Qian, M. F. Yin, Y. M. Zhuang, and K. Yu, “Cluster adaptive training for deep neural network,” in *Proc. 2015 IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, Brisbane, QLD, Australia, 2015, pp. 4325–4329.
- [97] C. Y. Wu and M. J. F. Gales, “Multi-basis adaptive neural network for rapid adaptation in speech recognition,” in *Proc. 2015 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Brisbane, QLD, Australia, 2015, pp. 4315–4319.
- [98] L. Samarakoon and K. C. Sim, “Factorized hidden layer adaptation for deep neural network based acoustic modeling,” *IEEE/ACM Trans. Audio Speech Language Processing*, vol. 24, no. 12, pp. 2241–2250, 2016.
- [99] L. Samarakoon, K. C. Sim, and B. Mak, “An investigation into learning effective speaker subspaces for robust unsupervised DNN adaptation,” in *Proc. 2017 IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, New Orleans, USA, 2017.
- [100] R. Kuhn, P. Nguyen, J. C. Junqua, L. Goldwasser, N. Niedzielski, S. Fincke, K. L. Field, and M. Contolini, “Eigenvoices for speaker adaptation,” in *Proc. 1998 IEEE the 5th Int. Conf. Spoken Language Processing*, Sydney, Australia, 1998, pp. 1774–1777.
- [101] M. J. F. Gales, “Cluster adaptive training for speech recognition,” in *Proc. 5th Int. Conf. Spoken Language Processing*, Sydney, Australia, 1998, pp. Article ID 0375.
- [102] M. Delcroix, K. Kinoshita, T. Hori, and T. Nakatani, “Context adaptive deep neural networks for fast acoustic model adaptation,” in *Proc. 2015 IEEE Int. Conf. Acoustics, Speech and Signal Processing* South Brisbane, QLD, Australia, 2015, pp. 4535–4539.
- [103] M. Delcroix, K. Kinoshita, C. Z. Yu, A. Ogawa, T. Yoshioka, and T. Nakatani, “Context adaptive deep neural networks for fast acoustic model adaptation in noisy conditions,” in *Proc. 2016 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Shanghai, China, 2016, pp. 5270–5274.
- [104] Y. Zhao, J. Li, K. Kumar, and Y. Gong, “Extended low-rank plus diagonal adaptation for deep and recurrent neural networks,” in *Proc. 2017 IEEE Int. Conf. Acoustics, Speech and Signal Processing*, New Orleans, USA, 2017.
- [105] M. Cooke, J. R. Hershey, and S. J. Rennie, “Monaural speech separation and recognition challenge,” *Computer Speech Lang.*, vol. 24, no. 1, pp. 1–15, Jan. 2010.
- [106] C. Weng, D. Yu, M. L. Seltzer, and J. Droppo, “Deep neural networks for single-channel multi-talker speech recognition,” *IEEE/ACM Trans. Audio Speech Lang Processing*, vol. 23, no. 10, pp. 1670–1679, Oct. 2015.
- [107] Y. X. Wang, A. Narayanan, and D. L. Wang, “On training targets for supervised speech separation,” *IEEE/ACM Trans. Audio Speech Lang Processing*, vol. 22, no. 12, pp. 1849–1858, Dec. 2014.
- [108] Y. Xu, J. Du, L. R. Dai, and C. H. Lee, “An experimental study on speech enhancement based on deep neural networks,” *IEEE Signal Processing Lett.*, vol. 21, no. 1, pp. 65–68, Jan. 2014.
- [109] F. Weninger, H. Erdogan, S. Watanabe, E. Vincent, J. Le Roux, J. R. Hershey, and B. Schuller, “Speech enhancement with LSTM recurrent neural networks and its application to noise-robust ASR,” in *Latent Variable Analysis and Signal Separation. LVA/ICA 2015. Lecture Notes in Computer Science*, Vincent E., Yeredor A., Koldovský Z., Tichavský P, Eds. Cham: Springer, 2015, pp. 91–99.
- [110] P. S. Huang, M. Kim, M. Hasegawa-Johnson, and P. Smaragdis, “Joint optimization of masks and deep recurrent neural networks for monaural source separation,” *IEEE/ACM Trans. Audio Speech Lang Processing*, vol. 23, no. 12, pp. 2136–2147, Dec. 2015.
- [111] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, “Deep clustering: Discriminative embeddings for segmentation and separation,” in *Proc. 2016 IEEE Int. Conf. Acoust. Speech Signal Process*, Shanghai, China, 2016, pp. 31–35.
- [112] Y. Isik, J. Le Roux, Z. Chen, S. Watanabe, and J. R. Hershey, “Single-channel multi-speaker separation using deep clustering,” in *17th Proc. Interspeech*, San Francisco, USA, 2016, pp. 545–549.
- [113] M. Cooke, *Modelling Auditory Processing and Organisation*. Cambridge: Cambridge Univ. Press, 2005.
- [114] D. P. Ellis, “Prediction-driven computational auditory scene analysis,” Ph.D. dissertation, Massachusetts: Massachusetts Inst. Technol., 1996.
- [115] M. Wertheimer, *Laws of organization in perceptual forms*, in *A Source Book of Gestalt Psychology*, W. D. Ellis, Ed. Trench: Trubner & Company, 1938.
- [116] M. N. Schmidt and R. K. Olsson, “Single-channel speech separation using sparse non-negative matrix factorization,” in *Proc. 2006-ICSLP the 9th Int. Conf. on Spoken Language Processing*, Pittsburgh, PA, USA, 2006.
- [117] P. Smaragdis, “Convolutional speech bases and their application to supervised speech separation,” *IEEE/ACM Trans. Audio Speech Lang. Processing*, vol. 15, no. 1, pp. 1–12, Jan. 2007.
- [118] J. Le Roux, F. Weninger, and J. Hershey, “Sparse NMF-half-baked or well done,” Mitsubishi Electr. Res. Labs (MERL), Cambridge, MA, USA, Tech. Rep. TR2015-023, Mar. 2015.
- [119] T. T. Kristjansson, J. R. Hershey, P. A. Olsen, S. J. Rennie, and R. A. Gopinath, “Super-human multi-talker speech recognition: the IBM 2006 speech separation challenge system,” in *Proc. 2006-ICSLP the 9th Int. Conf. Spoken Language Processing*, Pittsburgh, PA, USA, 2006, Article ID 1775-Mon1WeS.7.
- [120] T. Virtanen, “Speech recognition using factorial hidden markov models for separation in the feature space,” in *Proc. 2006-ICSLP 9th Int. Conf. Spoken Language Processing*, Pittsburgh, PA, USA, 2006.
- [121] R. J. Weiss and D. P. W. Ellis, “Monaural speech separation using source-adapted models,” in *Proc. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, New Paltz, NY, USA, 2007, pp. 114–117.
- [122] Z. Ghahramani and M. I. Jordan, “Factorial hidden Markov models,” *Mach. Learn.*, vol. 29, no. 2-3, pp. 245–273, Nov. 1997.
- [123] Z. Chen, Y. Luo, and N. Mesgarani, “Deep attractor network for single-microphone speaker separation,” in *Proc. 2017 IEEE Int. Conf. Acoust. Speech Signal Process*, New Orleans, USA, 2017.
- [124] D. Yu, X. Chang, and Y. M. Qian, “Recognizing multi-talker speech with permutation invariant training,” in *18th Proc. Interspeech*, Stockholm, Sweden, 2017.
- [125] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” in *Proc. the 27th Int. Conf. Neural Information Processing Systems*, Montreal, Canada, 2014, pp. 2672–2680.

- [126] Y. Shinohara, "Adversarial multi-task learning of deep neural networks for robust speech recognition," in *17th Proc. Interspeech*, San Francisco, USA, 2016, pp. 2369–2372.
- [127] D. Serdyuk, K. Audhkhasi, P. Brakel, B. Ramabhadran, S. Thomas, and Y. Bengio, "Invariant representations for noisy speech recognition," arXiv:1612.01928, 2016.
- [128] S. N. Sun, B. B. Zhang, L. Xie, and Y. N. Zhang, "An unsupervised deep domain adaptation approach for robust speech recognition," *Neurocomputing*, to be published.
- [129] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by backpropagation," arXiv:1409.7495, 2014.
- [130] R. Lippmann, E. Martin, and D. Paul, "Multi-style training for robust isolated-word speech recognition," in *Proc. IEEE Int. Conf. ICASSP '87 Acoustics, Speech, and Signal Processing*, Dallas, TX, USA, USA, 1987, pp. 705–708.
- [131] T. Ko, V. Peddinti, D. Povey, M. L. Seltzer, and S. Khudanpur, "A study on data augmentation of reverberant speech for robust speech recognition," in *Proc. 2017 IEEE Int. Conf. Acoust. Speech Signal Process*, New Orleans, USA, 2017.
- [132] J. Li, M. L. Seltzer, X. Wang, R. Zhao, and Y. Gong, "Large-scale domain adaptation via teacher-student learning," in *18th Proc. Interspeech*, 2017.
- [133] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," arXiv:1503.02531, 2015.
- [134] K. Markov and T. Matsui, "Robust speech recognition using generalized distillation framework," in *17th Proc. Interspeech*, San Francisco, USA, 2016, pp. 2364–2368.
- [135] S. Watanabe, T. Hori, J. Le Roux, and J. R. Hershey, "Student-teacher network learning with enhanced features," in *Proc. 2017 IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, Broadway, USA, 2017.
- [136] Z. Y. Lu, V. Sindhvani, and T. N. Sainath, "Learning compact recurrent neural networks," in *Proc. 2016 IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, Shanghai, China, 2016, pp. 5960–5964.
- [137] R. Prabhavalkar, O. Alsharif, A. Bruguier, and L. McGraw, "On the compression of recurrent neural networks with an application to LVCSR acoustic modeling for embedded speech recognition," in *Proc. 2016 IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, Shanghai, China, 2016, pp. 5970–5974.
- [138] W. Chan, N. R. Ke, and I. Lane, "Transferring knowledge from a RNN to a DNN," arXiv:1504.01483, 2015.
- [139] K. J. Geras, A. R. Mohamed, R. Caruana, G. Urban, S. J. Wang, Ö. Aslan, M. Philipose, M. Richardson, and C. Sutton, "Blending LSTMs into CNNs," arXiv:1511.06433, 2015.
- [140] L. Lu, M. Guo, and S. Renals, "Knowledge distillation for small-footprint highway networks," in *Proc. 2017 IEEE Int. Conf. Acoustics Speech and Signal Processing*, New Orleans, USA, 2017.
- [141] J. Cui, B. Kingsbury, B. Ramabhadran, G. Saon, T. Sercu, K. Audhkhasi, A. Sethy, M. Nussbaum-Thom, and A. Rosenberg, "Knowledge distillation across ensembles of multilingual models for low-resource languages," in *ICASSP*, 2017.
- [142] J. Li, R. Zhao, J. T. Huang, and Y. F. Gong, "Learning small-size DNN with output-distribution-based criteria," in *15th Proc. Interspeech*, Singapore, Singapore, 2014, pp. 1910–1914.
- [143] W. Xiong, J. Droppo, X. Huang, F. Seide, M. Seltzer, A. Stolcke, D. Yu, and G. Zweig, "Achieving human parity in conversational speech recognition," arXiv:1610.05256, 2016.
- [144] G. Saon, G. Kurata, T. Sercu, K. Audhkhasi, S. Thomas, D. Dimitriadis, X. D. Cui, B. Ramabhadran, M. Picheny, L. L. Lim, B. Roomi, and P. Hall, "English conversational telephone speech recognition by humans and machines," arXiv:1703.02136, 2017.
- [145] V. Vanhoucke, A. Senior, and M. Z. Mao, "Improving the speed of neural networks on CPUs," in *Proc. 2011 Deep Learning and Unsupervised Feature Learning NIPS Workshop*, Granada, Spain, 2011.
- [146] R. Alvarez, R. Prabhavalkar, and A. Bakhtin, "On the efficient representation and execution of deep acoustic models," arXiv:1607.04683v1, 2016.
- [147] R. Takeda, K. Nakadai, and K. Komatani, "Acoustic model training based on node-wise weight boundary model for fast and small-footprint deep neural networks," *Computer Speech & Language*, to be published.
- [148] Y. Q. Wang, J. Li, and Y. F. Gong, "Small-footprint high-performance deep neural network-based speech recognition using split-VQ," in *Proc. 2015 IEEE Int. Conf. Acoustics, Speech and Signal Processing Brisbane, QLD, Australia*, 2015, pp. 4984–4988.
- [149] V. Vanhoucke, M. Devin, and G. Heigold, "Multiframe deep neural networks for acoustic modeling," in *Proc. 2013 IEEE Int. Conf. Acoustics, Speech and Signal Processing Vancouver, BC, Canada*, 2013, pp. 7582–7585.
- [150] G. Pundak and T. N. Sainath, "Lower frame rate neural network acoustic models," in *17th Proc. Interspeech*, San Francisco, USA, 2016, pp. 22–26.



Dong Yu is a distinguished scientist and vice general manager at Tencent AI Lab. Prior to joining Tencent, he was a principal researcher at Microsoft Research. His research has been focusing on speech recognition and other applications of machine learning techniques. He has published two monographs and more than 160 papers in these areas and is the coinventor of 50+ granted and 10+ pending patents and the leader of CNTK. His works have been recognized by the prestigious IEEE Signal Processing Society 2013 and 2016 best paper awards and the ACMSE 2005 best paper award. Dr. Dong Yu is currently serving as a member of the IEEE Speech and Language Processing Technical Committee (2013–2018), the vice chair of IEEE Seattle section (2017–), and the distinguished lecturer of APSIPA (2017–2018). He has served as an associate editor of the *IEEE/ACM Transactions on Audio, Speech, and Language Processing* (2011–2015), an associate editor of the *IEEE Signal Processing Magazine* (2008–2011), the lead guest editor of the *IEEE Transactions on Audio, Speech, and Language Processing*—special issue on deep learning for speech and language processing (2010–2011), a guest editor of the *IEEE/CAA Journal of Automatica Sinica*—special issue on deep learning in audio, image, and text processing (2015–2016), and members of organization and technical committees of many conferences and workshops.



Jinyu Li received the bachelor and master degree from University of Science and Technology of China, in 1997 and 2000, with the highest honor, and the Ph.D. degree from Georgia Institute of Technology, Atlanta, in 2008. From 2000 to 2003, he was a researcher in the Intel China Research Center and Research Manager in iFlytek Speech, China. Currently, he is a principal applied scientist and technical lead in Microsoft Corporation, Redmond, WA. He leads a team to design and improve speech modeling algorithms and technologies that ensure industry state-of-the-art speech recognition accuracy for Microsoft products such as Cortana and xBox Kinect. His major research interests cover several topics in speech recognition, including deep learning, noise robustness, discriminative training, feature extraction, and machine learning methods. He authored more than 70 refereed publications and around 20 patents. He is the leading author of the book *Robust Automatic Speech Recognition—A Bridge to Practical Applications*, Academic Press, Oct, 2015. Currently, he serves as the associate editor of *IEEE/ACM Transactions on Audio, Speech and Language Processing*.