



End-to-end Adaptation with Backpropagation through WFST for On-device Speech Recognition System

Emiru Tsunoo, Yosuke Kashiwagi, Satoshi Asakawa, Toshiyuki Kumakura

Sony Corporation, Japan

Emiru.Tsunoo@sony.com, Yosuke.Kashiwagi@sony.com, Satoshi.Asakawa@sony.com,
Toshiyuki.Kumakura@sony.com

Abstract

An on-device DNN-HMM speech recognition system efficiently works with a limited vocabulary in the presence of a variety of predictable noise. In such a case, vocabulary and environment adaptation is highly effective. In this paper, we propose a novel method of end-to-end (E2E) adaptation, which adjusts not only an acoustic model (AM) but also a weighted finite-state transducer (WFST). We convert a pretrained WFST to a trainable neural network and adapt the system to target environments/vocabulary by E2E joint training with an AM. We replicate Viterbi decoding with forward-backward neural network computation, which is similar to recurrent neural networks (RNNs). By pooling output score sequences, a vocabulary posterior for each utterance is obtained and used for discriminative loss computation. Experiments using 2–10 hours of English/Japanese adaptation datasets indicate that the fine-tuning of only WFSTs and that of only AMs are both comparable to a state-of-the-art adaptation method, and E2E joint training of the two components achieves the best recognition performance. We also adapt each language system to the other language using the adaptation data, and the results show that the proposed method also works well for language adaptations.

Index Terms: speech recognition, end-to-end training, weighted finite-state transducer, environment adaptation, vocabulary adaptation, language adaptation

1. Introduction

Conventional automatic speech recognition (ASR) systems, typically referred to as DNN-HMM ASRs, commonly consist of DNN acoustic models (AMs), which estimate the posterior probabilities of hidden Markov model (HMM) states discriminatively, and language models (LMs), which are typically word n -grams [1]. Triphone HMMs and LMs are trained separately and incorporated together into a weighted finite-state transducer (WFST) for computation efficiency. Further, in the case of on-device ASR systems, both AMs and LMs are compressed to reduce computational cost [2, 3]. However, framewise cross-entropy training of AMs does not always minimize word error rates (WERs) because LMs model word sequences that have different granularities. To fill the gap between the AMs and LMs, sequence discriminative training has been studied and has significantly reduced WERs [4–8].

Recently, end-to-end (E2E) ASRs have attracted much attention as methods of directly integrating of AMs and LMs, including connectionist temporal classification (CTC) [9–11], attention-based encoder-decoder models [12–14], and hybrid models [15, 16]. As the E2E ASRs directly map speech input frames into an output label sequence, the parameters are expected to be more optimized than those in conventional DNN-HMM ASRs in which modules are trained separately. In general, E2E ASRs, which are superior to conventional DNN-

HMM ASRs, require thousands of hours of speech-transcription parallel data [17, 18].

In the cases of ASR systems such as those in a robot, where most of the noise sounds are predictable, environment adaptation is highly effective. Further, if the vocabulary is limited, better performance is expected using the adaptation data of such vocabulary. Among the adaptation methods, there are three major directions, input feature transformation [19, 20], the use of auxiliary features [21, 22], and the direct adaptation or transformation of DNN parameters [23–26], in which the last is related to this work the most. However, all of them are concerned only with acoustic modeling.

In this paper, we propose a novel method of E2E adaptation. We convert a pretrained WFST to a trainable neural network and adapt the system to the target environments/vocabulary by E2E joint retraining with the pretrained AM. This conversion is invertible; thus, the trained WFST parameters can be converted into an ordinary WFST and utilized in a conventional manner. We replicate Viterbi decoding with forward-backward neural network computation, which is similar to recurrent neural networks (RNNs). For an application of on-device command recognition with a vocabulary size of hundreds of words, a vocabulary posterior for each utterance is obtained by simply pooling output score sequences, which is used for discriminative loss computation. Experiments using 2–10 hours of English/Japanese adaptation data indicate that the fine-tuning of only the WFSTs and that of only the AMs are both comparable to a state-of-the-art adaptation method, and E2E joint training of the two components achieves the best recognition performance. We also adapt each language system to the other language using the adaptation data, and the results show that the proposed method also works well for language adaptations.

The rest of the paper is organized as follows. In Section 2 we describe our proposed WFST conversion algorithm to a trainable neural network. In Section 3 we briefly give some guidelines for an efficient implementation. We describe our experiments in Section 4 and conclude the paper in Section 5.

2. Backpropagation through WFST

To exploit pretrained AMs and LMs, and to improve the performance with a small amount of adaptation data, we convert a WFST to make it jointly trainable. Although sequence discriminative training optimizes the AM training in a joint criterion with a given WFST [4–8], to the best of our knowledge, no previous study has been carried out on updating a WFST during the training process, because it is complicated in general and the update might interfere with the LM. Since Viterbi decoding is a simple linear procedure, it is differentiable. Therefore, we alternate Viterbi decoding with a combination of neural network modules, and we see a DNN AM and the neural network decoder based on a WFST as a joint E2E system that is to be




Figure 1: ViterbiNet: Neural network structure for Viterbi decoding.

optimized. The structure of the neural network is shown in Figure 1 and we refer to this network as ViterbiNet in this paper.

2.1. Forward-backward Computation

First, each AM posterior is mapped to WFST states associated with the posterior by a sparse affine transformation. This is equivalent to the transition model mapping followed by ilabel (input id of arc) selection for each state in Kaldi terminology. Denoting $\mathbf{x}^{(t)} \in \mathbb{R}^P$ as the P -dimensional posterior of time frame t (an output of the AM neural network), the transformation can be described as

$$\tilde{\mathbf{x}}^{(t)} = \mathbf{W}\mathbf{x}^{(t)}, \quad (1)$$

where each element of $\mathbf{W} \in \mathbb{R}^{S \times P}$ is defined as

$$w_{i,j} = \begin{cases} 1 & (\text{TransitionIdToPdf}(j) = \text{ilabel}(i)) \\ 0 & (\text{otherwise}) \end{cases} \quad (2)$$

and S is the total state number. The matrix \mathbf{W} is constant and kept unchanged in this paper.

The WFST transition can be represented as a sparse matrix $\mathbf{V} \in \mathbb{R}^{S \times S}$. An example of a WFST is illustrated in Figure 2, where the triphone HMM transitions and self-loops are omitted for simplicity. In this case, the graph is represented as \mathbf{V} in Figure 2. Let the forward score for each state at time frame t be $\boldsymbol{\alpha}^{(t)} \in \mathbb{R}^S$. This can be computed with matrix \mathbf{V} as

$$\begin{aligned} \boldsymbol{\alpha}^{(t)} &= f(\tilde{\mathbf{x}}^{(t)}, \boldsymbol{\alpha}^{(t-1)}) \\ &= \mathbf{V} \otimes \left(\boldsymbol{\alpha}^{(t-1)} \circ \tilde{\mathbf{x}}^{(t)} \right), \end{aligned} \quad (3)$$

where \circ is element-wise multiplication and \otimes is the maximum value in each row after multiplying the elements of \mathbf{V} by those of $\boldsymbol{\alpha}^{(t-1)} \circ \tilde{\mathbf{x}}^{(t)}$. Thus, $\boldsymbol{\alpha}^{(t)}$ can also be written as

$$\alpha_i^{(t)} = \max_j v_{i,j} \alpha_j^{(t-1)} \tilde{x}_j^{(t)}. \quad (4)$$

Typically, the initial state of a WFST is set to 0; therefore, the initial score $\boldsymbol{\alpha}^{(0)}$ is

$$\alpha_i^{(0)} = \begin{cases} 1 & (i = 0) \\ 0 & (\text{otherwise}). \end{cases} \quad (5)$$

The above manipulation is mathematically equivalent to Viterbi decoding, and this can be carried out recursively, similarly to RNNs.

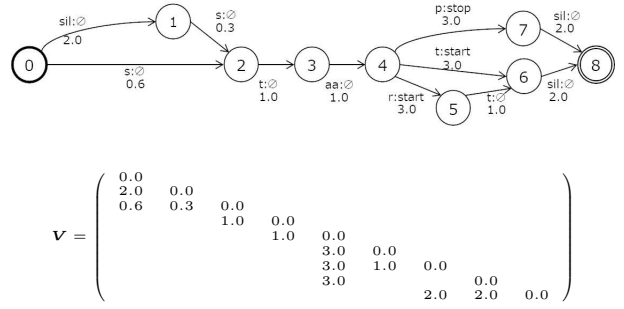


Figure 2: Example of sparse matrix representation of WFST.

In general, to decode the output sequence, all olabels (output ids of arcs) of highest-scoring forward paths have to be memorized. Instead, in this paper, we also compute the backward score, similarly to (3), as

$$\boldsymbol{\beta}^{(t)} = \left(\mathbf{V}^T \otimes \boldsymbol{\beta}^{(t+1)} \right) \circ \tilde{\mathbf{x}}^{(t+1)}. \quad (6)$$

The final score of the WFST, $\boldsymbol{\beta}^{(T)}$, can be defined using the final score of each state. By combining forward and backward computations, the scores of output symbols are calculated, which is described in the following section.

2.2. Output Layer

Each arc between states has an olabel, and if the olabel is not $\langle \text{eps} \rangle$, then an output symbol is emitted at each time frame. Therefore, the output score $\mathbf{y}^{(t)} \in \mathbb{R}^C$ with vocabulary size C is defined as

$$y_u^{(t)} = \max_{v_{i,j} \in \mathbb{U}} \alpha_i^{(t-1)} v_{i,j} \tilde{x}_i^{(t)} \beta_j^{(t)}, \quad (7)$$

where \mathbb{U} is a set of arcs whose olabels are u . This resembles to Baum-Welch algorithm, except it takes maximum instead of summation for decoding. Then, the score sequence \mathbf{y} can be aligned with the reference word sequence using a similar method to the forward-backward algorithm in CTC training [27]. However, in this paper, only a small vocabulary size is used; therefore, it is feasible to concatenate all words of each command and represent them as single words (e.g., ‘‘Take-a-picture’’). Taking this into account, the score of the output command u can be simply written as

$$\hat{y}_u = \max_t y_u^{(t)}, \quad (8)$$

which is realized with maxpooling manipulation.

2.3. Loss Function

The decoding is simplified to a command classification using the output score $\hat{\mathbf{y}}$. The loss function is

$$\begin{aligned} L &= \text{CrossEntropy}\left(\frac{\hat{\mathbf{y}}}{\|\hat{\mathbf{y}}\|_1}, u\right) \\ &= \text{SoftMaxCrossEntropy}(\log \hat{\mathbf{y}}, u), \end{aligned} \quad (9)$$

where $\|\cdot\|_1$ is the l_1 -norm of the vector. Since all aforementioned transformations including affine, forward-backward, and output computations are linear manipulations, the gradients for backpropagation are straightforwardly derived.

If only the AM is updated with the gradients, it acts similarly to sequence discriminative training [4–8]. The E2E training is carried out by updating both the AM and the matrix \mathbf{V} , which is the WFST itself. In the case of a simplified single-word model, the loss function becomes exactly equivalent to the MMI criterion [4].

3. Implementation

We have theoretically discussed the proposed method. To realize it in feasible manner, several guidelines are given.

3.1. WFST Preprocess

The WFST parameters are converted to a sparse matrix \mathbf{V} and applied in every time step. In general, however, the WFST contains a number of transitions that do not require any emissions, i.e., epsilon arcs. As long as the WFST contains such nonemitting transitions, the proposed forward–backward procedure cannot be realized with the aforementioned recursive scheme. Therefore, the WFST should first be processed with `fstrmepsilon` to remove epsilon arcs as much as possible. In many cases, not all the epsilons are removed but the number becomes much smaller (e.g., 0.036% of arcs for Switchboard trigram LM). The surviving epsilon arcs are omitted from the matrix \mathbf{V} for training and only utilized for evaluation in this paper.

3.2. Sparse Computation

The mapping transform matrix \mathbf{W} in (1) and the transition matrix \mathbf{V} in (3) are sparse matrices. The use of ordinary affine transform modules is impractical, especially for large WFST parameters. Therefore, we encode the sparse values in the matrices and compute only these elements. This is inefficient and memory-consuming when carried out using GPUs. We leave this computation to the CPU.

3.3. Log-scale Computation

The proposed forward–backward computation also requires a series of multiplications over the input utterance. This leads to rounding errors in computation. Although normalization of the values in each step can improve the precision, we adopt log-scale computation for the recursive multiplication similarly to Viterbi decoding. In this case, the multiplications in (3) and (6) turn into summations of logarithms, and the gradients of the forward–backward layer become constant values delivered from the output layer and spread along the possible transitions. Therefore, the log-scale computation increases the speed of inference as well as backpropagation.

3.4. Regularization of AM

Sequence discriminative training tends to overfit [7] and the proposed E2E training is not an exception. Since we exploit the pretrained AM, we adopt Kullback–Leibler (KL) divergence regularization, as in [24]. By adding this divergence as a regularization term to (9), the optimization criterion becomes

$$L' = \text{SoftMaxCrossEntropy}(\log \hat{\mathbf{y}}, \hat{\mathbf{r}}) + \lambda \sum_{t=1}^T \text{KL}(\mathbf{x}_{\text{org}}^{(t)} || \mathbf{x}^{(t)}), \quad (10)$$

where $\mathbf{x}_{\text{org}}^{(t)}$ is the AM posterior of time frame t computed with the original pretrained parameters, and λ is the controllable regularization weight. The regularization allows not only the AM to remain close to the original AM, but also the fast convergence

Table 1: *Adaptation and evaluation datasets for command recognition tasks*

Task	EN-SC	EN-ROBOT	JP-ROBOT
examples	“cat” “three”	“come-here” “move-forward”	“oide” “zenshin”
Adaptation data			
vocabulary	20	157	229
# of utterances	34,760 (9.6 h)	3,297 (1.8 h)	4,580 (2.1 h)
# of speakers	1,034	7	20
SNR	0 – 20 dB	-5 – 15 dB	-5 – 15 dB
Evaluation data			
vocabulary	20	157	229
# of utterances	29,961 (8.3 h)	1,413 (0.7 h)	2,290 (1.0 h)
# of speakers	847	3	10
SNR	0 – 20 dB	-5 – 15 dB	-5 – 15 dB

because it spreads peaky gradients of the ViterbiNet over the dimensions of the posterior that represent similar HMM states. The regularization of WFST parameters is left as future work.

4. Experiments

We evaluated command recognition tasks with small-footprint setups for on-device systems such as those in robots. US English and Japanese systems were built for the evaluation. Three tasks were evaluated; English Speech Commands (EN-SC, 20 commands) [28], English robot commands (EN-ROBOT, 157 commands), and Japanese robot commands (JP-ROBOT, 229 commands). The proposed method was compared with state-of-the-art sequence discriminative training and adaptation methods. We also conducted language adaptations from Japanese to English and from English to Japanese using the same task setups.

4.1. Experimental Setup

The input speech signals were 16 kHz, 16-bit monaural speech data processed by multi-scale multi-band DenseNet-based noise reduction (NR) described in [29], which was trained using WSJ, TIMIT, and ATR-503 Japanese data. The feature vector for every 10 ms was a 24-dimensional log filterbank and energy with Δ and $\Delta\Delta$ (75 dimensions in total). The features were spliced into 11 frames (5 on each side of the current frame) followed by a 5-layer full connection DNN with 640 hidden units for each layer. We built two AM DNNs, for English and Japanese, where the outputs of the AM DNNs were 1,256 units for the English model, and 1,012 units for the Japanese model. The AMs were first trained solely with 28,000 hours of English data and 12,300 hours of Japanese data, respectively, which were mixed with mechanical and living noises, followed by fine-tuning with a 280 hour English subset and a 165 hour Japanese subset jointly with the front-end DenseNet NR. We used these AMs for the baseline systems.

LMs were built for each task. For the EN-SC task, a grammar-based WFST with a vocabulary of 20 words was built, which had 493 states and 1,452 arcs. For the EN-ROBOT task, we constructed a grammar-based WFST of 225 commands, including all commands in the adaptation and evaluation sets (8,351 states, 21,879 arcs), and a grammar WFST of 305 commands (3,797 states, 8,014 arcs), including all adaptation/evaluation commands for the JP-ROBOT task.

For adaptation, we converted the WFSTs to ViterbiNets and jointly retrained them with the AMs using adaptation datasets (ViterbiNet E2E), specifically, by updating all parameters in the 5-layer AM DNN and the matrix \mathbf{V} . The proposed network

Table 2: Sentence error rates (SERs) of adaptation methods in command recognition tasks

Adaptation methods	EN-SC	EN-ROBOT	JP-ROBOT
(No adaptation)	34.70	8.99	6.24
Cross Entropy	10.78	6.44	4.06
sMBR [30]	10.83	6.65	3.14
KL-regularization [24]	10.77	6.23	3.28
ViterbiNet AM	9.65	6.09	2.93
ViterbiNet WFST	13.08	3.89	3.06
ViterbiNet E2E	9.26	3.54	2.66

was implemented with the open source *Sony Neural Network Libraries*¹. Training was carried out using the Adam optimizer (learning rate = 0.0001, $\beta_1 = 0.9$, $\beta_2 = 0.999$) with a 16-utterance minibatch, and 20 epochs were sufficient for convergence in every task. The regularization parameter λ in (10) was set to 0.01. We also evaluated the cases where only the AMs were updated (ViterbiNet AM) in the proposed scheme, which is theoretically equivalent to sequence discriminative training, and only the ViterbiNets, i.e., the matrix \mathbf{V} , were updated (ViterbiNet WFST). We also compared the above retraining with the simple retraining of only AMs using aligned labels (Cross Entropy), sMBR discriminative training [30], and state-of-the-art KL-regularization adaptation [24], using the same adaptation dataset with 40 iterations for sMBR and 20 epochs for the other methods. The adaptation datasets contained all commands in the evaluation sets, but all speakers were different; thus only the environment and vocabulary matched.

In the evaluation step, trained ViterbiNet parameters were converted back to WFSTs, and Viterbi decoding was applied as in conventional ASR. The beam size was set to 7 and the acoustic scale was 0.07. To all the EN-SC adaptation data and evaluation data, we added randomly chosen noise in the dataset (e.g., running tap) with a uniform distribution between an SNR of 0 dB and 20 dB before processing with the fine-tuned DenseNet NR. To all the adaptation and evaluation datasets for EN-ROBOT and JP-ROBOT, we randomly added mechanical and living noise (e.g., servomotor sounds, music) an SNR from -5 dB to 15 dB before the NR. The details of the adaptation and evaluation datasets are given in Table 1.

4.2. Environment/Vocabulary Adaptation Result

The baselines and all the obtained sentence error rates (SERs) of adaptation methods are shown in Table 2. For sMBR, the lowest SERs among the iterations are shown. Compared with the baselines, all adaptations performed better since the conditions of vocabulary and environments matched. The EN-SC task performed badly because the AM was trained with completely different noise mixtures, and all the commands in the task were simple and short; therefore, it was difficult to distinguish them. Moreover, the environments varied since the data was collected from a number of volunteers in various rooms with different microphones. The results of the ViterbiNet AM were similar to those of sMBR or even better since they were equivalent to the conventional MMI training, not in the state level but in the word level. Surprisingly, updating only ViterbiNet WFST also significantly reduced SERs. We suppose that the trained WFST can model sequential information rather than static information, which worked well, especially for the vocabulary adaptations. The joint E2E training of the AM and

¹Neural Network Libraries by Sony: <https://nnabla.org/>

Table 3: SERs of adaptation methods in command recognition tasks using AMs of different languages as initial seeds

Adaptation methods	JP AM→	JP AM→	EN AM→
	EN-SC	EN-ROBOT	JP-ROBOT
(No adaptation)	84.12	94.13	46.03
KL-regularization [24]	67.35	54.64	16.94
ViterbiNet AM	54.40	31.21	10.39
ViterbiNet WFST	47.24	49.54	25.68
ViterbiNet E2E	27.64	13.09	7.64

WFST using the proposed ViterbiNet consistently achieved the best performances.

4.3. Language Adaptation

Next, we evaluated language adaptation from Japanese to English using the EN-SC and EN-ROBOT tasks and from Japanese to English using the JP-ROBOT task. For the English tasks we used the Japanese AM as a seed, and new LMs were compiled to match the AM using Japanese phones and HMM. The lexicons were designed manually, such as *seven* for “seven” and *three* for “three,” since Japanese does not have *v* and *th* pronunciations. Similarly, we used the English AM as a seed for the Japanese task and compiled the LM using English phones and HMM. Then we applied each method to adapt to the language of the test tasks using the aforementioned adaptation datasets. The results are shown in Table 3. The original results in the first row were poor because of the language mismatches of AMs and the ill-designed LMs. However, with adaptations, we obtained similar tendencies to the results in Table 2 and achieved the best performance using the proposed E2E adaptation. The E2E adaptation was so strong that it approached the performances of language-matched training, and the results imply that a few hours of adaptation data are sufficient to build a system for a task as simple as command recognition tasks, even from a model of a different language. We also tested the adaptations from random initial values for the AMs, but the training did not progress at all.

5. Conclusions and Future Work

We have proposed a new end-to-end adaptation method by fine-tuning an AM as well as a WFST. We converted the pretrained WFST to a trainable neural network and adapted the system to target environments/vocabulary. We replicated Viterbi decoding with forward-backward neural network computation, which is similar to RNNs. By pooling the output score sequences, a vocabulary posterior for each utterance was obtained and used for a discriminative loss computation. Experiments using 2–10 hours of English/Japanese adaptation data indicated that the fine-tuning of only WFSTs and that of only AMs were both comparable to a state-of-the-art adaptation method, and the E2E joint training of the two components achieved the best recognition performances. We also adapted one language system to the other language using a small amount of adaptation data, and the results showed that the proposed method also worked well for language adaptations.

Our ongoing work includes training larger WFSTs, such as Switchboard trigram LM, for LVCSR. Also, the proposed training method can be applied to external LMs for E2E ASR systems using CTC/Attention, where the complexity should be lower and the training becomes more feasible, especially with character/subword outputs.

6. References

- [1] H. Bourlard and N. Morgan, *Connectionist Speech Recognition*. Kluwer Academic Publishers, 1994.
- [2] X. Lei, A. W. Senior, A. Gruenstein, and J. Sorensen, "Accurate and compact large vocabulary speech recognition on mobile devices," in *Proc. of Interspeech*, 2013, pp. 662–665.
- [3] Y. Wang, J. Li, and Y. Gong, "Small-footprint high-performance deep neural network-based speech recognition using split-VQ," in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference*, 2015, pp. 4984–4988.
- [4] D. Povey, "Discriminative training for large vocabulary speech recognition," Ph.D. dissertation, University of Cambridge, 2005.
- [5] B. Kingsbury, "Lattice-based optimization of sequence classification criteria for neural-network acoustic modeling," in *Acoustics, Speech and Signal Processing (ICASSP), 2009 IEEE International Conference*, 2009, pp. 3761–3764.
- [6] K. Vesely, A. Ghoshal, L. Burget, and D. Povey, "Sequence-discriminative training of deep neural networks," in *Proc. of Interspeech*, 2013, pp. 2345–2349.
- [7] H. Su, G. Li, D. Yu, and F. Seide, "Error back propagation for sequence training of context-dependent deep networks for conversational speech transcription," in *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference*, 2013, pp. 6664–6668.
- [8] D. Povey, V. Peddinti, D. Galvez, P. Ghahremani, V. Manohar, X. Na, Y. Wang, and S. Khudanpur, "Purely sequence-trained neural networks for ASR based on lattice-free MMI," in *Proc. of Interspeech*, 2016, pp. 2751–2755.
- [9] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of the 23rd International Conference on Machine Learning*. ACM, 2006, pp. 369–376.
- [10] A. Graves and N. Jaitly, "Towards end-to-end speech recognition with recurrent neural networks," in *International Conference on Machine Learning*, 2014, pp. 1764–1772.
- [11] Y. Miao, M. Gowayyed, and F. Metze, "EESSEN: End-to-end speech recognition using deep RNN models and WFST-based decoding," in *Automatic Speech Recognition and Understanding (ASRU), 2015 IEEE Workshop*, 2015, pp. 167–174.
- [12] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, "Attention-based models for speech recognition," in *Advances in Neural Information Processing Systems*, 2015, pp. 577–585.
- [13] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference*, 2016, pp. 4960–4964.
- [14] L. Lu, X. Zhang, and S. Renais, "On training the recurrent neural network encoder-decoder for large vocabulary end-to-end speech recognition," in *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference*, 2016, pp. 5060–5064.
- [15] S. Kim, T. Hori, and S. Watanabe, "Joint CTC-attention based end-to-end speech recognition using multi-task learning," in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference*, 2017, pp. 4835–4839.
- [16] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, "Hybrid CTC/attention architecture for end-to-end speech recognition," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [17] R. Prabhavalkar, K. Rao, T. N. Sainath, B. Li, L. Johnson, and N. Jaitly, "A comparison of sequence-to-sequence models for speech recognition," in *Proc. of Interspeech*, 2017, pp. 939–943.
- [18] J. Li, G. Ye, R. Zhao, J. Droppo, and Y. Gong, "Acoustic-to-word model without OOV," in *Automatic Speech Recognition and Understanding Workshop (ASRU), 2017 IEEE*, 2017, pp. 111–117.
- [19] J. Neto, L. Almeida, M. Hochberg, C. Martins, L. Nunes, S. Renais, and T. Robinson, "Speaker-adaptation for hybrid HMM-ANN continuous speech recognition system," in *Fourth European Conference on Speech Communication and Technology*, 1995.
- [20] K. Yao, D. Yu, F. Seide, H. Su, L. Deng, and Y. Gong, "Adaptation of context-dependent deep neural networks for automatic speech recognition," in *Spoken Language Technology Workshop (SLT), 2012 IEEE*, 2012, pp. 366–369.
- [21] G. Saon, H. Soltan, D. Nahamoo, and M. Picheny, "Speaker adaptation of neural network acoustic models using i-vectors," in *ASRU*, 2013, pp. 55–59.
- [22] M. Delcroix, K. Kinoshita, C. Yu, A. Ogawa, T. Yoshioka, and T. Nakatani, "Context adaptive deep neural networks for fast acoustic model adaptation in noisy conditions," in *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference*, 2016, pp. 5270–5274.
- [23] R. Gemello, F. Mana, S. Scanzio, P. Laface, and R. De Mori, "Adaptation of hybrid ANN/HMM models using linear hidden transformations and conservative training," in *Acoustics, Speech and Signal Processing (ICASSP), 2006 IEEE International Conference*, vol. 1, 2006, pp. 1189–1192.
- [24] D. Yu, K. Yao, H. Su, G. Li, and F. Seide, "KL-divergence regularized deep neural network adaptation for improved large vocabulary speech recognition," in *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference*, 2013, pp. 7893–7897.
- [25] P. Swietojanski and S. Renais, "Learning hidden unit contributions for unsupervised speaker adaptation of neural network acoustic models," in *Spoken Language Technology Workshop (SLT), 2014 IEEE*, 2014, pp. 171–176.
- [26] O. Klejch, J. Fainberg, and P. Bell, "Learning to adapt: A meta-learning approach for speaker adaptation," in *Proc. of Interspeech*, 2018, pp. 867–871.
- [27] A. Graves, "Sequence transduction with recurrent neural networks," in *ICML Representation Learning Workshop*, 2012.
- [28] P. Warden, "Speech commands: A public dataset for single-word speech recognition." *Dataset available from http://download.tensorflow.org/data/speech_commands_v0.01.tar.gz*, 2017.
- [29] N. Takahashi and Y. Mitsufuji, "Multi-scale multi-band densenets for audio source separation," in *Applications of Signal Processing to Audio and Acoustics (WASPAA), 2017 IEEE Workshop*, 2017, pp. 21–25.
- [30] M. Gibson and T. Hain, "Hypothesis spaces for minimum Bayes risk training in large vocabulary speech recognition," in *Ninth International Conference on Spoken Language Processing*, 2006.