

LEARNING DISCRIMINATIVE FEATURES IN SEQUENCE TRAINING WITHOUT REQUIRING FRAMEWISE LABELLED DATA

Jun Wang¹, Dan Su¹, Jie Chen¹, Shulin Feng³, Dongpeng Ma¹, Na Li¹, Dong Yu²

¹Tencent AI Lab, Shenzhen, China ²Tencent AI Lab, Bellevue WA, USA

³Key Laboratory on Machine Perception, Peking University, Beijing, China

{joinerwang, dansu, leojiechen, shulinfeng, dongpengma, nenali, dyu}@tencent.com

ABSTRACT

In this work, we try to answer two questions: Can deeply learned features with discriminative power benefit an ASR system's robustness to acoustic variability? And how to learn them without requiring frame-wise labelled sequence training data? As existing methods usually require knowing where the labels occur in the input sequence, they have so far been limited to many real-world sequence learning tasks. We propose a novel method which simultaneously models both the sequence discriminative training and the feature discriminative learning within a single network architecture, so that it can learn discriminative deep features in sequence training that obviates the need for presegmented training data. Our experiment in a realistic industrial ASR task shows that, without requiring any specific fine-tuning or additional complexity, our proposed models have consistently outperformed state-of-the-art models and significantly reduced Word Error Rate (WER) under all test conditions, and especially with highest improvements under unseen noise conditions, by relative 12.94%, 8.66% and 5.80%, showing our proposed models can generalize better to acoustic variability.

Index Terms— ASR, discriminative feature learning, acoustic variability, sequence discriminative training.

1. INTRODUCTION

Mismatch between training and testing conditions is a ubiquitous problem in modern neural network-based systems. It is particularly common in perceptual sequence learning tasks (e.g. speech recognition, handwriting recognition, motion recognition) where unsegmented noisy input streams are annotated with strings of discrete labels. For example, acoustic variability for modern Automatic Speech Recognition (ASR) systems arises from speaker, accent, background noise, reverberation, channel, and recording conditions, etc. Ensuring robustness to variability is a challenge, as it is impractical to see all types of acoustic conditions during training.

Several model and feature based adaptation methods has been proposed, such as Noise Adaptive Training [1, 2, 3] and

Vector Taylor Series [4] for handling environment variability, and Maximum Likelihood Linear Regression (MLLR) [5] and iVectors [6] for handling speaker variability. Discriminative feature learning is another important orientation. For example, manifold regularized deep neural networks [7] has been proposed for speech-in-noise recognition tasks. While manifold learning approaches are known to require relatively high complexity, recently, alternative methods for learning discriminative features have been originally suggested for Face Recognition (FR) tasks. Among them, Contrastive Loss [8, 9] and Triplet Loss [10] respectively construct loss functions for image pairs and triplets to supervise the embedding learning; Center Loss [11] encourages intra-class compactness by penalizing the distances between the features of samples and their centers; Alternative approaches such as SphereFace [12] and CosFace [13] replace the simple Euclidean distance metric with more suitable metric space for FR tasks. Variants of these methods have also been successfully adopted in Speaker Recognition (SR) tasks [14, 15].

For FR and SR tasks, their goal of learning discriminative features is to identify new unseen classes without label prediction[11, 16]. In contrast, our goal for ASR tasks is to enhance robustness to new unseen acoustic conditions during testing. In Section 2.1, we first investigate the deeply learned features with a joint supervision of frame-level Cross Entropy (CE) Loss and Center Loss in our ASR system. Despite the success of discriminative feature learning approaches for FR and SR problems, so far, it has not been possible to apply these approaches to sequence-discriminative training models. The problem is that these loss functions are defined separately for each point in the training sequence; even worse, any other loss functions adopted for the joint supervision must also be defined separately for each point, thus sequence-discriminative training mechanisms, such as Connectionist Temporal Classification (CTC)[17], MMI[18], MBR[19] and LFMMI[20]), are ruled out. To address this problem, in Section 2.2, we propose a novel, general mechanism by encouraging the sequence-discriminative training model to simultaneously learn feature representations with discriminative power, so that the model can exploit the full

³Shulin did this work while he was an intern at Tencent

potential of sequence modelling mechanisms. Section 3 compares our proposed mechanisms to baseline methods, and the paper concludes with Section 4.

2. PROPOSED METHOD

During testing, an ASR system may encounter new recording conditions, microphone types, speakers, accents and background noises. Even if the test scenarios are seen during training, there can be significant variability in their statistics. Can deeply learned representations with more discriminative power benefit the system’s robustness to acoustic variability? In Section 2.1, we firstly consider fusion of two criterias for framewise training, and then in Section 2.2 we discuss two losses for joint supervision for sequence training without requiring framewise labels and investigate their potential effects on feature learning.

2.1. Framewise Multi-loss Fusion

We assume an acoustic model whose output nodes represent K classes, e.g., context dependent (CD) phonemes, CD sub-phonemes, or labels of the associated HMM states, among others. Center Loss (CL) can be formulated as follows:

$$\mathcal{L}_{cl} = \sum_t \|\mathbf{u}_t - \mathbf{c}_{k_t}\|_2^2 \quad (1)$$

where the deeply learned feature $\mathbf{u}_t \in \mathbb{R}^D$ is the second last layer’s output at time t ; given framewise training data and annotation pairs $(x_t, k_t) : t = 1, \dots, T$, the input x_t belongs to the k_t -th class, and the $\mathbf{c}_{k_t} \in \mathbb{R}^D$ denotes the k_t -th class center of the deep features.

Center Loss learns a center of deep features for each class and penalizes the distance between the deep features and their corresponding class centers; it can significantly enhance the discriminative power of deep features [11, 16]. However, without separability, the deep features and centers could degrade to zero at which point the Center Loss is very small. Therefore we use a Cross Entropy (CE) Loss function to ensure the separability of the deep features:

$$\mathcal{L}_{ce} = - \sum_t \log y_t^{k_t} \quad (2)$$

where $y_t^{k_t}$ is the k_t -th output after the softmax operation on the output of the last layer:

$$y_t^{k_t} = \frac{e^{\mathbf{a}_t^{k_t}}}{\sum_{j=1}^K e^{\mathbf{a}_t^j}} \quad (3)$$

$$\mathbf{a}_t = W\mathbf{u}_t + B \quad (4)$$

where $\mathbf{a}_t \in \mathbb{R}^K$ is the output of the last layer and \mathbf{a}_t^j indicates the j -th element, and $W \in \mathbb{R}^{K \times D}$, $B \in \mathbb{R}^K$ are the weight and bias matrix of the last layer.

The Framewise Multi-Loss Fusion (FMF) of the above two losses with a balancing scalar λ can then be formulated as follows:

$$\mathcal{L}_{fmf} = \mathcal{L}_{ce} + \lambda\mathcal{L}_{cl} \quad (5)$$

The joint supervision by FMF benefits both the inter-class separability and the intra-class compactness of deep features, and will be proved in our experiment to be effective for robustness to acoustic variability.

2.2. Temporal Multi-loss Fusion

This section presents a novel method called Temporal Multi-loss Fusion (TMF) for learning discriminative deep features that removes the need for framewise labelled training sequences. It models both the sequence discriminative training and the feature discriminative learning within a single network architecture. The basic idea is to interpret the network outputs as a probability distribution over all possible label sequences, conditioned on a given sequence. Given this distribution, a fusion loss function can be derived that directly maximizes the probabilities of the correct labelings while penalizing the distance between the deep features and the corresponding centers.

Without loss of generality, we illustrate sequence discriminative training with CTC [17] method. The goal of Maximum Likelihood (ML) training in CTC is to simultaneously maximize the log probabilities of all the correct classifications in the training set. This means minimizing the following loss function:

$$\mathcal{L}_{ml} = - \sum_{(\mathbf{x}, \mathbf{z}) \in S} \ln(p(\mathbf{z}|\mathbf{x})) \quad (6)$$

where $(\mathbf{x}, \mathbf{z}) \in S$ are the training data pairs. Since the framewise annotated training data $(x_t, k_t) : t = 1, \dots, T$ is no longer available, we define a conditional Expected Center Loss (ECL) as follows:

$$\mathcal{L}_{ecl} = \sum_s \sum_t p(s, t|\mathbf{z}) \|\mathbf{u}_t - \mathbf{c}_{z'_s}\|_2^2 \quad (7)$$

where \mathcal{L}_{ecl} means a conditional expected loss by the deep representation \mathbf{u}_t deviating from $\mathbf{c}_{z'_s}$ which is the center corresponding to symbol z'_s . Following the convention of CTC, for labelling sequence \mathbf{z} of length r , we denote by $\mathbf{z}_{1:s}$ and $\mathbf{z}_{r-s:r}$ its first and last s symbols respectively, define a set of positions where label j occurs as $lab(\mathbf{z}, j) = \{s : \mathbf{z}'_s = j\}$, where a modified label sequence \mathbf{z}' is defined with blanks added to the beginning and the end of \mathbf{z} , and inserted between every pair of consecutive labels. Given the labelling \mathbf{z} , the probability of all the paths corresponding to \mathbf{z} that go through a given symbol s in \mathbf{z}' at time t can be calculated as the product of the forward and backward variables at the symbol s and time t :

$$p(s, t|\mathbf{z}) = \alpha_t(s)\beta_t(s) \quad (8)$$

where the detailed calculation for the forward variables $\alpha_t(s)$ and backward variables $\beta_t(s)$ can be referred to [17].

Finally, we formulate the TMF objective function as below:

$$\mathcal{L}_{tmf} = \mathcal{L}_{ml} + \lambda \mathcal{L}_{ecl} \quad (9)$$

where a scalar λ is used for balancing ECL and CTC loss. ECL encourages the intra-class compactness while CTC Loss encourages the separability of features. Consequently, their joint supervision minimizes the intra-class variations while keeping the features of different classes separable.

As the objective function is differentiable, the network can then be trained with standard back-propagation through time. We provide the learning details about TMF in Algorithm 1, and accordingly illustrate its feedforward process and backpropagate error signal flow in Figure 1.

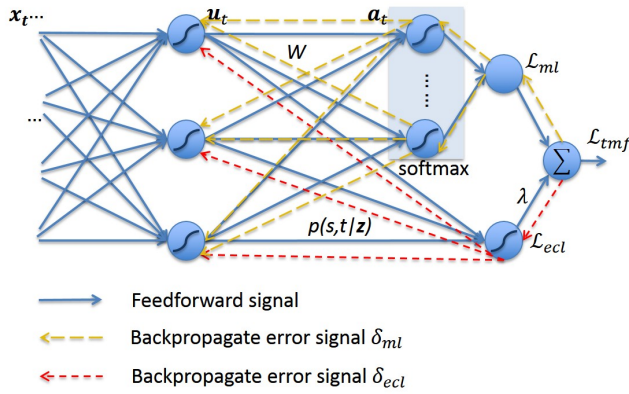


Fig. 1. Illustration of the feedforward process and the backpropagate error signal flow under fusion supervision: The feedforward signal is denoted with blue solid lines; the backpropagate error signals generated by CTC and ECL are denoted with yellow and red dash lines, respectively.

3. EXPERIMENTS AND ANALYSIS

3.1. Data

A noisy far-field training corpus was simulated based on a clean corpus of 2 million clean speech utterances which were 1845.78 hours in total, by adding reverberation and mixing with various environmental noises. The room simulator was based on IMAGE method [21] and generated 15K room impulse responses (RIRs). The noisy speech corpus covered reverberation time (RT60) ranging from 0 to 600ms, and covered different noise types at uniformly distributed SNRs ranging from 15 to 30 dB. Both the clean and the noisy corpus were split into training set of 1840 hours and validation set of 5.78 hours which contains 6K utterances.

To construct test set, we used a separated clean speech dataset of real-world conversation recordings, denoted as

Algorithm 1 Temporal Multi-loss Fusion learning algorithm

Input: Training pairs of sequences $(\mathbf{x}, \mathbf{z}) \in S$; initialized parameters θ in convolutional and LSTM layers, parameters W and $\{\mathbf{c}_j | j = 1, 2, \dots, K\}$ in full connection and loss layers, respectively; balancing factor λ , and learning rate μ and γ .

Output: The parameters θ and W .

- 1: **while** not converge **do**
- 2: Compute the TMF joint loss by Equation 9
- 3: Compute the backpropagation error signal through the softmax layer:

$$\delta_{ml}(k) = \frac{\partial \mathcal{L}_{tmf}}{\partial \mathbf{a}_t^k} = y_t^k - \frac{\sum_{s \in \text{lab}(z,k)} \alpha_t(s) \beta_t(s)}{\sum_{s=1}^{|z'|} \alpha_t(s) \beta_t(s)} \quad (10)$$

- 4: Compute the backpropagation error signal by ECL:

$$\delta_{ecl} = \frac{\partial \mathcal{L}_{ecl}}{\partial \mathbf{u}_t} = \sum_s \alpha_t(s) \beta_t(s) (\mathbf{u}_t - \mathbf{c}_{z'_s}) \quad (11)$$

- 5: Compute the fusion of backpropagation error signals through the second last layer:

$$\delta = W^T \delta_{ml} + \lambda \delta_{ecl} \quad (12)$$

- 6: Compute weight adjustments ΔW and $\Delta \theta$ with the chain rule using the backpropagation error signals δ_{ml} and δ , respectively.
- 7: Update the parameters:

$$\mathbf{c}_{z'_s} = \mathbf{c}_{z'_s} - \mu \sum_t \alpha_t(s) \beta_t(s) (\mathbf{c}_{z'_s} - \mathbf{u}_t) \quad (13)$$

$$W = W - \gamma \Delta W \quad (14)$$

$$\theta = \theta - \gamma \Delta \theta \quad (15)$$

- 8: **end while**
-

test_clean. It had 1K utterances and 1.43 hours length in total. We simulated a noisy test set by mixing *test_clean* with the same noise types as training set and denoted it as *noise_seen*. Furthermore, to evaluate the robustness of our proposed method under unseen acoustic conditions, we simulated another noisy test set by mixing *test_clean* with a different set of noise types not seen during training and denoted this test set as *noise_unseen*. Both *noise_seen* and *noise_unseen* were simulated with RT60 ranging from 0 to 600ms, with uniformly distributed SNRs from 15 to 30dB.

FBank features of 40-dimension were computed with 25ms window length, 10ms hop size, and formed a 120-dimension vector along with their first and second order difference. After normalization, the vector of the current frame was concatenated with those of the preceding 5 frames and the subsequent 5 frames, resulting in an input vector of dimension $40 * 3 * (5 + 1 + 5) = 1320$.

3.2. Experiment Setup

While increasing the neural network depth, as discussed in [3], the internal deep representations became increasingly discriminative and insensitive to many sources of variability in speech signals, although DNNs could not extrapolate to test samples substantially different from the training examples which were not sufficiently representative. Hence it was important for us to set a solid bar as baseline systems. We adopted state-of-the-art models which have been well-tuned for our industrial applications as our baselines: For frame-wise training, *CE* indicated the model trained with CE loss function as our baseline model; correspondingly, *FMF* indicated the proposed model with framewise multi-loss joint supervision. For sequence training, *CTC* indicated the model trained with CTC criteria as our baseline model, and *TMF* indicated the proposed model with temporal multi-loss joint supervision; furthermore, we wanted to look into the potential benefit by the proposed *TMF* model to other state-of-the-art ASR system: The model denoted as *CTC+MMI* was initialized with the *CTC* model and then trained with sequence discriminative training [22, 23] with MMI criteria [18]; the model denoted as *TMF+MMI* was initialized with the *TMF* model and then trained with MMI.

All the above models shared the same basal network architecture and hyperparameter configuration, so that they had approximately equal computation complexity, as the complexity added by joint supervision could be negligible. The architecture contained two 2-dimensional CNN layers, each with a kernel size of (3, 3), a stride of (1, 1), and followed with a maxpool layer with a kernel size of (2, 2) and a stride of (2, 2), and then five LSTM layers, each with hidden size of 1024 and with peephole, and then one full-connection FC layer plus a softmax layer. Batch normalization was applied after each CNN and LSTM layer to accelerate convergence and improve generalization. We used CD phonemes as our output units, which were about 12K classes in our Chinese ASR system.

During training, the balancing factor λ was set to exponential values of $1e-3$ and $1e-4$ for clean and noisy conditions, respectively. The class centers were updated only when $\alpha_t(s)\beta_t(s) \geq 0.01$, with the batch momentum μ of $1e-3$, and the "blank" class was excluded. Adam optimizer was adopted. The learning rate had an initial value of $1e-4$ and would be halved if the average validation likelihood after every 5K batches hadn't raised for 3 successive times. The training would be early stopped if the likelihood hadn't raised for 8 successive times.

3.3. Result and Analysis

Table 1 compares experiment result in terms of test WERs by different models: *CE* versus *FMF* for framewise training, *CTC* versus *TMF*, and *CTC+MMI* versus *TMF+MMI* for sequence training. The better performances are marked with bold num-

bers; their corresponding relative improvements are given in parentheses, and the highest relative improvements among *test_clean*, *noise_seen* and *noise_unseen* test conditions are marked with underlines.

It shows that the proposed models consistently outperform their corresponding baseline models under all test conditions. Meanwhile, it is worth noting that all of the highest relative improvements are achieved under the unseen noise condition. This result proves that the joint supervision by our proposed methods effectively benefit the system's robustness, and can generalize the sequence discriminative training better to noise variability by enhancing the inter-class separability and the intra-class compactness of deep feature.

We obtained the above improvements without fine-tuning hyperparameters for the proposed models. As mentioned in Section 3.2, the network structure and hyperparameters have been fine-tuned to optimize the baseline models, and then directly inherited by the proposed models. Hence another merit of the proposed models is that they don't necessarily require additional complexity or hyperparameter tuning in addition to a given baseline model.

Table 1. WERs (%) of *CE* v.s. *FMF*, *CTC* v.s. *TMF*, and *CTC+MMI* v.s. *TMF+MMI*; the better scores were marked with bold numbers, followed with their relative improvement in parentheses; the highest relative improvement scores among various noise conditions were underlined.

| | <i>test_clean</i> | <i>noise_seen</i> | <i>noise_unseen</i> |
|----------------|--------------------|--------------------|------------------------------|
| <i>CE</i> | 5.66 | 8.82 | 10.05 |
| <i>FMF</i> | 5.19 (8.30) | 8.15 (7.60) | 8.75 (<u>12.94</u>) |
| <i>CTC</i> | 4.67 | 8.19 | 8.08 |
| <i>TMF</i> | 4.33 (7.28) | 7.59 (7.33) | 7.38 (<u>8.66</u>) |
| <i>CTC+MMI</i> | 4.45 | 7.34 | 7.58 |
| <i>TMF+MMI</i> | 4.29 (3.59) | 7.12 (3.00) | 7.14 (<u>5.8</u>) |

4. CONCLUSIONS

We have introduced a novel method which obviates the need for framewise labelled training data and allows the network to directly learn deep discriminative feature representations while performing sequence discriminative training. We have demonstrated that the deeply learned features with discriminative power could benefit the ASR system's robustness to noise variability. Without requiring any specific fine-tuning or additional complexity, the proposed models have outperformed state-of-the-art models for an industrial ASR system, and the relative improvements are especially remarkable under unseen conditions. Future research includes extending the proposed method to various sequence discriminative training mechanisms, and evaluating their effectiveness to various sources of acoustic variability.

5. REFERENCES

- [1] O. Kalinli, M.L. Seltzer, J. Droppo, and A. Acero, “Noise adaptive training for robust automatic speech recognition,” *IEEE Transactions on Audio, Speech, and Language Processing*, pp. 1889–1901, 2010.
- [2] M.L. Seltzer, D. Yu, and Y. Wang, “An investigation of deep neural networks for noise robust speech recognition,” in *Proceedings of International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2013.
- [3] D. Yu, M.L. Seltzer, J. Li, J.T. Huang, and F. Seide, “Feature learning in deep neural networks - studies on speech recognition tasks,” in *Proceedings of International Conference on Learning Representation*. IEEE, 2013.
- [4] C.K. Un and N.S. et al. Kim, “Speech recognition in noisy environments using first-order vector Taylor series,” *Speech Communication*, vol. 24(1), pp. 39–49, 1998.
- [5] Y. Wang and M.J. Gales, “Speaker and noise factorization for robust speech recognition,” *IEEE Transactions on Audio, Speech and Language Processing*, vol. 20(7), pp. 2149–2158, 2012.
- [6] G. Saon, H. Soltau, D. Nahamoo, and M. Picheny, “Speaker adaptation of neural network acoustic models using i-vectors,” in *Proceedings of IEEE Automatic Speech Recognition and Understanding Workshop*. IEEE, 2013, pp. 55–59.
- [7] V. S. Tomar and R. Rose, “Manifold regularized deep neural networks,” in *Proceedings of Conference of the International Speech Communication Association*. IEEE, 2014.
- [8] R. Hadsell, S. Chopra, and Y. LeCun, “Dimensionality reduction by learning an invariant mapping,” in *Proceedings of Computer vision and pattern recognition (CVPR)*. IEEE computer society conference, 2006, vol. 2, pp. 1735–1742.
- [9] Y. Sun, Y. Chen, X. Wang, and X. Tang, “Deep learning face representation by joint identification-verification,” in *Advances in Neural Information Processing Systems*, 2014, pp. 1988–1996.
- [10] F. Schro, D. Kalenichenko, and J. Philbin, “Facenet: A unified embedding for face recognition and clustering,” in *Proceedings of CVPR*, 2015, pp. 815–823.
- [11] Y.D. Wen, K.P. Zhang, Z.F. Li, and Y. Qiao, “A discriminative feature learning approach for deep face learning,” in *Proceedings of European Conference on Computer Vision*, 2016.
- [12] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song, “Sphereface: Deep hypersphere embedding for face recognition,” in *Proceedings of CVPR*. IEEE, 2017, pp. 212–220.
- [13] H. Wang, Y.T. Wang, Z. Zhou, X. Ji, D.H. Gong, J.C. Zhou, Z.F. Li, and W. Liu, “Cosface: Large margin cosine loss for deep face recognition,” in *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2018.
- [14] S. Yadav and A. Rai, “Learning discriminative features for speaker identification and verification,” in *Proceedings of Conference of the International Speech Communication Association*. IEEE, 2018.
- [15] N. Li, D.Y. Tuo, D. Su, Z.F. Li, and D. Yu, “Deep discriminative embeddings for duration robust speaker verification,” in *INTERSPEECH*. IEEE, 2018.
- [16] H. Liu, W. Shi, W.P. Huang, and Q. Guan, “A discriminatively learned feature embedding based on multi-loss fusion for person search,” in *Proceedings of ICASSP*. IEEE, 2018.
- [17] A. Graves, A. Fernandez, 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*. IEEE, 2006.
- [18] D. Povey, D. Kanevsky, B. Kingsbury, B. Ramabhadran, G. Saon, and K. Visweswariah, “Boosted mmi for model and feature-space discriminative training,” in *Proceedings of ICASSP*. IEEE, 2008.
- [19] B. Kingsbury, “Lattice-based optimization of sequence classification criteria for neural network acoustic modeling,” in *Proceedings of ICASSP*. IEEE, 2009.
- [20] 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 *INTERSPEECH*. IEEE, 2016.
- [21] J.B. Allen and D.A. Berkley, “Image method for efficiently simulation room-small acoustic,” *Journal of the Acoustical Society of America*, vol. 65, pp. 943–950, 1979.
- [22] K. Vesely, A. Ghoshal, L. Burget, and D. Povey, “Sequence-discriminative training of deep neural networks,” in *INTERSPEECH*, 2013.
- [23] H. Sak, O. Vinyals, G. Heigold, A. Senior, E. McDermott, R. Monga, and M. Mao, “Sequence discriminative distributed training of long short-term memory recurrent neural networks,” in *INTERSPEECH*, 2014.