



# Self-Teaching Networks

Liang Lu, Eric Sun and Yifan Gong

Microsoft

{liang.lu, sun.eric, yifan.gong}@microsoft.com

## Abstract

We propose self-teaching networks to improve the generalization capacity of deep neural networks. The idea is to generate soft supervision labels using the output layer for training the lower layers of the network. During the network training, we seek an auxiliary loss that drives the lower layer to mimic the behavior of the output layer. The connection between the two network layers through the auxiliary loss can help the gradient flow, which works similar to the residual networks. Furthermore, the auxiliary loss also works as a regularizer, which improves the generalization capacity of the network. We evaluated the self-teaching network with deep recurrent neural networks on speech recognition tasks, where we trained the acoustic model using 30 thousand hours of data. We tested the acoustic model using data collected from 4 scenarios. We show that the self-teaching network can achieve consistent improvements and outperform existing methods such as label smoothing and confidence penalization.

**Index Terms:** speech recognition, deep neural network, self-teaching regularizer

## 1. Introduction

Deep neural networks trained with millions of parameters have achieved tremendous success in various machine learning tasks including speech, natural language and image processing [1]. Despite the success, there are still some challenges to be addressed in deep learning. For example, although trained with huge amount of data, the neural networks still cannot generalize well to the scenarios which are not well represented by the training data, and overfit to domains which dominate the training set. One example is speech recognition in the ‘wild’, where a single speech recognizer is deployed in various acoustic conditions and domains including far- or near-field and different dialects. Due to the vanishing gradient problem, another challenge is to train deeper networks that can fully harness the information from the vast amount of data. Residual and highway networks [2, 3] have made significant progress along this direction, however, their applications in recurrent networks are not as successful as in feedforward and convolutional networks. Training very deep networks with high generalization capacity is still a challenging optimization problem in the research community.

There have been significant research efforts devoted to improve the generalization ability of the deep neural networks. Traditional approaches include the use of weight noise, L1/L2 regularization, and dropout [4]. However, these approaches are not widely used in training commercial speech recognition systems where the amount of training data is in thousands to tens of thousands hours, as their effects tend to diminish when the amount of data increases. Recently, a few regularization approaches for neural networks have been proposed that introduce noise to the supervision labels, e.g., label smoothing [5] and confidence penalization [6]. The rationale behind these ap-

proaches is that in the supervised learning task, the ground truth labels are from a Multinoulli distribution with the probability of the correct class being 1 and others being 0. However, for neural networks which use softmax as the output layer, the predictions are soft probabilities, which means they cannot completely match the ground truth labels. Training the network with more epochs can always push down the cross-entropy loss, but with the cost of overfitting to the training set, indicating a lower generalization capacity. To preventing overfitting and improve generalization, a random noise is applied to the labels in label smoothing, allowing the model to match an ‘easier’ target distribution. Motivated in a similar way, confidence penalization aims at penalizing the entropy of the output distribution of the neural network to avoid the network placing the entire probability to a single class.

In this paper, we propose self-teaching networks – another approach to tackle the vanishing gradient problem in deep networks as well as improving the generalization capacity of the network. In contrast to residual and highway networks, which feed the features from the lower layers into the upper layers with skip connections, in self-teaching networks, we pass the supervision signal from the top layer down to the lower layer directly. To this end, we use the top layer to generate the supervision labels, and use them as the additional supervision signal to train the lower layers. This is analogous to teacher-student training [7, 8, 9], where the teacher model generates the soft labels to train the student model. But the key difference here is that the ‘teacher’ and ‘student’ now refer to different layers of the *same* network. Thus we name this approach as ‘self-teaching’. Given the soft supervision labels from the output layer, we seek an auxiliary loss to minimize the error compared to the soft labels. While it benefits the gradient flow due to the skip connection, the auxiliary loss also works as a regularizer, which improves the generalization capacity of the network. While this is a general training approach applicable to different types of network architectures, in this study, we narrow down the scope and focus on deep recurrent networks with long short-term memory (LSTM) units [10]. In terms of applications, we evaluate the self-teaching networks on speech recognition tasks, where we trained the acoustic model using 30 thousand hours of data. Our test dataset consists of audio recorded from 4 different scenarios, and we show consistent improvements by the self-teaching network as well as better generalization compared to label smoothing and confidence penalization.

## 2. Method

When the neural network goes deeper, the gradients to the lower layers become noisy and less informative. An indication of this behavior is the underfitting phenomena as observed in [2]. In fact, Zhang et al. [11] shows that the lower layers of the network may be more important to extract the useful feature representations in classification tasks. Residual and highway networks [2, 3] tackle this problem by introducing skip connections

which feed the lower level features directly to the upper layers. In this case, the lower layers will be closer to the supervision signals and they can receive higher quality gradient signals for parameter update.

Motivated in the similar way, we also draw connections between the top and lower layers in self-teaching networks, but it works the other way round. Rather than directly feeding the lower layer features upward through a skip connection, we pass down the supervision signal from the output layer to the lower layers. More precisely, suppose the posterior probability from the output layer of the network is  $P_\theta(\mathbf{y}|\mathbf{h}^L)$ , where  $\mathbf{y}$  denotes the output,  $\mathbf{h}^L$  is the activation from the top layer of the network,  $L$  denotes the network depth and  $\theta$  is the set of model parameters. We then add another softmax layer to a lower layer of the network to compute the posterior probability  $P_\theta(\mathbf{y}|\mathbf{h}^l)$ , where  $1 \leq l < L$ . During the network training, we also minimize the distance between the two distributions together with the original cross-entropy loss function in order to encourage the lower layer to follow the behaviour of the top layer. If we use the Kullback-Leibler (KL) divergence as the measure between the two posterior distributions, the loss function becomes

$$\begin{aligned} \mathcal{L}(\theta) &= H(\bar{\mathbf{y}}, P_\theta(\mathbf{y}|\mathbf{h}^L)) + \lambda D_{KL}(P_\theta(\mathbf{y}|\mathbf{h}^L) || P_\theta(\mathbf{y}|\mathbf{h}^l)) \\ &= H(\bar{\mathbf{y}}, P_\theta(\mathbf{y}|\mathbf{h}^L)) + \lambda H(P_\theta(\mathbf{y}|\mathbf{h}^L), P_\theta(\mathbf{y}|\mathbf{h}^l)) \\ &\quad - \lambda H(P_\theta(\mathbf{y}|\mathbf{h}^L)) \end{aligned} \quad (1)$$

where  $\bar{\mathbf{y}} \in \{0, 1\}^Z$ ,  $\sum_j \bar{y}_j = 1$  is the ground truth label with the number of classes as  $Z$ ;  $\lambda$  is the tunable parameter;  $H(\mathbf{p})$  is the entropy of the distribution  $\mathbf{p}$ ,

$$H(\mathbf{p}) = - \sum_i \mathbf{p}_i \log \mathbf{p}_i, \quad (2)$$

and  $H(\mathbf{p}, \mathbf{q})$  denotes the cross-entropy of the two distributions  $\mathbf{p}$  and  $\mathbf{q}$ , i.e.,

$$H(\mathbf{p}, \mathbf{q}) = - \sum_i \mathbf{p}_i \log \mathbf{q}_i \quad (3)$$

for discrete distributions.  $D_{KL}(\mathbf{p}||\mathbf{q})$  denotes the KL-divergence between the two distributions, i.e.,

$$D_{KL}(\mathbf{p}, \mathbf{q}) = H(\mathbf{p}, \mathbf{q}) - H(\mathbf{p}). \quad (4)$$

An example of the self-teaching network with LSTMs is shown in Figure 1, which we will focus on in this paper.

To compute the probability  $P_\theta(\mathbf{y}|\mathbf{h}^l)$ , it requires an additional linear projection layer to compute the logits (unnormalized probabilities) before the softmax operation. If the number of classes  $Z$  is large, this can result in considerable increase in terms of the model size during training, although this layer is discarded during inference, which means there is no additional computational cost to deploy the model. When the amount of training data is relatively small, it may be reasonable to share the linear projection layer with the output layer to reduce the total number of model parameters in training. However, this is not investigated as we have abundant training data as detailed in the experimental section.

For speech recognition, the posterior probabilities  $P_\theta(\mathbf{y}|\mathbf{h}^L)$  and  $P_\theta(\mathbf{y}|\mathbf{h}^l)$  can be at the frame-level or at the sequence-level. In the frame-level scenario, the posterior probabilities are obtained simply by feeding in the acoustic frames to the network followed by the softmax operation.

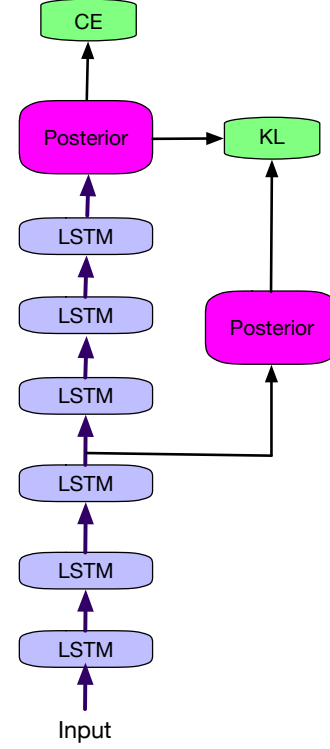


Figure 1: An example of self-teaching network with a 6-layer LSTM studied in this work. If the posterior is at the frame-level, we use cross-entropy (CE) as the main loss function. The auxiliary connection is from the top layer to the 3rd layer of the network.

The loss function is in the form of cross-entropy. The loss function corresponding to the self-teaching network in this case is shown by Eq (1). For the case of sequence-level scenario, the two probabilities can be obtained by forward-backward computation over the lattices, and the main cross-entropy loss function  $H(\bar{\mathbf{y}}, P_\theta(\mathbf{y}|\mathbf{h}^L))$  should be replaced by the one corresponding to maximum mutual information (MMI) or minimum Bayesian risk (MBR), while the KL-divergence term remains the same. In this work, we only focus on the frame-level scenario, and leave the sequence-level training as our future work.

### 3. Related Work

As mentioned above, the self-teaching network proposed in this paper is related to several previous works in the literature. In this section, we only briefly review a few closely related ones.

#### 3.1. Label smoothing

Label smoothing [5] is a simple method to add noise to the supervision labels to increase the robustness of the network. It can be viewed as sampling from a noise distribution  $\mathbf{q}$  and adding the output to the ground truth labels  $\bar{\mathbf{y}}$ . The noise perturbed labels are then renormalized to follow a valid distribution, which are then used as the new supervision labels. As shown in [5], it is equivalent to minimize the follow loss function:

$$\mathcal{L}(\theta) = H(\bar{\mathbf{y}}, P_\theta(\mathbf{y}|\mathbf{h}^L)) + \lambda D_{KL}(\mathbf{q} || P_\theta(\mathbf{y}|\mathbf{h}^L)). \quad (5)$$

The noise distribution  $\mathbf{q}$  can be uniform as used in [5, 12, 13], or a Multinoulli distribution as in [14]. Compared to the self-teaching network as Eq (1), the key difference is that the noise distribution is replaced by the posterior computed from a lower layer output of the network.

### 3.2. Confidence penalization

Confidence penalization [6] aims at penalizing the output distribution of the network which places entire probability to a single class. The method is to introduce a penalization term that encourages the output distribution to have a higher entropy. From the information theory, a distribution reaches its maximum entropy when it is uniform. Consequently, the effect of the penalization term is to push the model to place reasonable probabilities to the entire classes, so that it will be closer to a uniform distribution. This is in effect similar to label smoothing with the uniform label noise. Concretely, the loss function of confidence penalization can be expressed as

$$\mathcal{L}(\theta) = H(\bar{\mathbf{y}}, P_{\theta}(\mathbf{y}|\mathbf{h}^L)) - \lambda H(P_{\theta}(\mathbf{y}|\mathbf{h}^L)). \quad (6)$$

Compared to Eq (1), it is the same as the loss function of the self-teaching network, but without the cross-entropy term between  $P_{\theta}(\mathbf{y}|\mathbf{h}^L)$  and  $P_{\bar{\theta}}(\mathbf{y}|\mathbf{h}^L)$ .

### 3.3. Teacher-Student learning

The self-teaching network is also related to the teacher-student learning [8, 9] approach for transfer learning between different networks. Let  $\bar{\theta}$  be the set of parameters of the teacher model, the loss function for teacher-student learning may be represented as

$$\mathcal{L}(\theta) = H(\bar{\mathbf{y}}, P_{\theta}(\mathbf{y}|\mathbf{h}^L)) + \lambda D_{KL}(P_{\bar{\theta}}(\mathbf{y}|\mathbf{h}^L) || P_{\theta}(\mathbf{y}|\mathbf{h}^L)). \quad (7)$$

Since the teacher model  $\bar{\theta}$  is not updated during the teacher-student training, the entropy term  $H(P_{\bar{\theta}}(\mathbf{y}|\mathbf{h}^L))$  is ignored during the gradient computation. Hence the loss can be simplified as

$$\mathcal{L}(\theta) = H(\bar{\mathbf{y}}, P_{\theta}(\mathbf{y}|\mathbf{h}^L)) + \lambda H(P_{\bar{\theta}}(\mathbf{y}|\mathbf{h}^L), P_{\theta}(\mathbf{y}|\mathbf{h}^L)). \quad (8)$$

In practice, the cross entropy from the ground truth labels  $H(\bar{\mathbf{y}}, P_{\theta}(\mathbf{y}|\mathbf{h}^L))$  is also ignored as it does not make much difference in both convergence and accuracy [15]. Without this term, it also allows the unlabeled data to be used in teacher-student training, which is a particular advantage. Comparing Eq (7) to Eq (1), the formulation of the self-teaching network is the same as that of the teacher-student learning, while the major difference is that both the teacher and student model are different components or layers from the same model in the self-teaching network.

### 3.4. Residual and highway networks

The self-teaching network shares some similarities with residual and highway networks [2, 3] in the sense that they all aim at tackling the vanishing gradient problem when propagating back the gradients to the lower layers of the network. They all create skip connections so that the lower layers are closer to the supervision signals. However, for self-teaching networks, the skip connection is to pass down the supervision signal from the

Table 1: *Statistics of the 4 evaluation sets used in our experiments. Cortana contains both close-talk and far-field speech. DMA-I refers to the in-domain distant microphone array condition, and DMA-O is the out-domain distant microphone array condition, which is not covered in the training data.*

Domain	# utterances	# words
Cortana	79105	438970
DMA-O	5485	29480
DMA-I	75056	284092

output layer, while for residual and highway networks, the skip connection is used to feed the lower layer features directly to the upper layers.

## 4. Experiments and Results

### 4.1. Experimental setup

We evaluated the self-teaching networks for speech recognition tasks. We focus on LSTMs for acoustic modeling as they are still the backbone for most of the state-of-the-art systems. In our experiments, the models were trained with 30 thousand hours of anonymized and transcribed Microsoft production data, recorded in both close-talk and far-field conditions. We work on hybrid systems, where the hidden Markov models (HMM) is used for sequence modeling. The number of tied triphone states is 9404 in our experiments, and the input feature is 80-dimension log Mel filter banks sampled every 20 milliseconds. The language model is a 5-gram with around 100 million (M) n-grams.

Given the vast amount of training data, it is very expensive to thoroughly evaluate the two key aspects of the self-teaching network, namely, its abilities to improve the generalization capacity and to help gradient flow for very deep networks. In this study, we mainly focus on the generalization aspect of the self-teaching network. Therefore, we used 6-layer LSTMs for all our experiments as shown in Figure 1, and only compare to the other two regularization approaches, i.e., label smoothing and confidence penalization. All the LSTM models were trained with truncated back-propagation through time (BPTT) with truncation length of 16. The acoustic models were evaluated on 4 test sets, which are detailed in Table 1. In particular, the DMA-O condition is not covered by the training data, which is primarily for evaluating the generalization power of the model.

### 4.2. Results of smaller models

Table 2 shows the results of a smaller 6-layer LSTM model in terms of word error rates (WERs), where the size of the cell is 1024 and the dimension of the hidden vector is 512 after a linear projection layer. The number of the total model parameters is roughly 31 million. We trained the model using the CNTK deep learning toolkit [16], and it took around 7 – 10 days with parallel training using 16 GPUs for each model to converge. We searched the hyper-parameter  $\lambda \in \{0.001, 0.005, 0.01, 0.02\}$ , and for clarity, we only show results with the optimal  $\lambda$ . In fact, we observed a consistent trend by tuning the hyper-parameter  $\lambda$ . While label smoothing has been shown to improve the accuracy for end-to-end models such as the encoder-decoder attention model in [12], and Connectionist Temporal Classification (CTC) model in [13], we did not observe consistent

Table 2: Experimental results of the smaller LSTM model with around 31 million parameters. LS refers to label smoothing, and CP is short for confidence penalization. ST refers to the self-teaching network. H means the entropy term  $H(P_\theta(\mathbf{y}|\mathbf{h}^L))$  in Eq (1).

System	WER		
	Cortana	DMA-O	DMA-I
Baseline	9.46	19.56	13.37
LS	9.56	19.35	13.27
CP	9.40	19.44	13.53
ST with H	9.32	19.13	13.29
ST without H	9.35	19.37	13.18

improvements for our hybrid model by varying the value of  $\lambda$ . Confidence penalization shows the similar trend. For self-teaching networks, however, we observed small but consistent improvements across the 4 evaluation sets. The new model is also fairly robust to the hyper-parameter in our experiments, as the results only fluctuated slightly when  $\lambda$  is within the range from 0.001 to 0.02. Since the difference between confidence penalty and self-teaching network is only the cross entropy term  $H(P_\theta(\mathbf{y}|\mathbf{h}^L), P_\theta(\mathbf{y}|\mathbf{h}^L))$ , we evaluated its impact by removing the entropy term  $H(P_\theta(\mathbf{y}|\mathbf{h}^L))$  in the experiments of self-teaching networks. As shown in Table 2, we can still observe consistent improvements without the entropy term, indicating that the cross-entropy term may be more important. This may help explain the results of confidence penalization, which show that the entropy term alone cannot achieve consistent WER reductions.

#### 4.3. Results of larger models

We then increased the size of the model with the size of cell being 1800 and the dimension of the hidden vector being 600. The total number of model parameters becomes roughly 63 million, which is about twice as many as the model has in the previous experiments. This model takes around 2 weeks to converge when trained with 16 GPUs in parallel. Due to the computational cost, we only run experiments with self-teaching networks, and the results are shown in Table 3. The large model improved the baseline system, with WER reductions vary from 3.8% for the Cortana task to 4% - 8% for the DMA task. For self-teaching networks, compared to the results of the smaller model, we observed relatively larger improvements after increasing the model size with the relative WER reduction up to 3% - 4%, which indicates that the regularization term plays a larger role in the model training. Notably, the improvement for the Cortana task is comparable to that achieved by doubling the model size. In fact, measured by the amount of training data in our experiments, the model size is not considerably large. We may achieve an even larger gain by increasing the model size further. Finally, we also evaluated the impact of the entropy term in these experiments, and the observation is consistent with that in the previous experiments.

### 5. Future Work

So far, we have investigated the self-teaching network with frame-level training for speech recognition. In the future, we shall study the case of sequence training, where the posterior probabilities are computed at the sequence-level by running the forward-backward algorithm over lattices. Furthermore, in this

Table 3: Experimental results of the larger LSTM model with roughly 63 million parameters.

System	WER		
	Cortana	DMA-O	DMA-I
Baseline	9.10	17.92	12.84
ST with H	8.84	17.32	12.70
ST without H	8.79	17.39	12.70

work, we only focus on the generalization aspect of the self-teaching network. In the future, we shall also evaluate this approach to train very deep networks, particularly for recurrent networks such as LSTMs. We have shown that the self-teaching network works better for larger models given the fixed amount of training data, which indicates it may achieve larger improvements in the case of lower resource conditions or languages, which will be studied in the future as well. The application of the self-teaching network for other tasks including speaker recognition is also an interesting research direction.

### 6. Conclusions

We proposed the self-teaching network to tackle the gradient vanishing problem in deep networks, and to improve the generalization capacity of the network. We approached the problem by introducing the skip connection between the output layer to a lower layer of the network to pass down the supervision signals. We investigated the application of the self-teaching network for speech recognition, where we trained the LSTM based acoustic models using 30 thousand hours of transcribed data. Compared to the closely related ideas such as label smoothing and confidence penalization, we have shown that self-teaching networks can achieve higher and consistent improvements across 4 evaluation datasets, and larger gains were observed with a larger LSTM model.

### 7. References

- [1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, 2015.
- [2] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *arXiv preprint arXiv:1512.03385*, 2015.
- [3] R. K. Srivastava, K. Greff, and J. Schmidhuber, "Training very deep networks," in *Proc. NIPS*, 2015.
- [4] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [5] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [6] G. Pereyra, G. Tucker, J. Chorowski, Ł. Kaiser, and G. Hinton, "Regularizing neural networks by penalizing confident output distributions." 2017. [Online]. Available: <https://arxiv.org/abs/1701.06548>
- [7] C. Bucilua, R. Caruana, and A. Niculescu-Mizil, "Model compression," in *Proc. ACM SIGKDD*, 2006.
- [8] J. Li, R. Zhao, J.-T. Huang, and Y. Gong, "Learning small-size DNN with output-distribution-based criteria," in *Proc. INTERSPEECH*, 2014.
- [9] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," in *Proc. NIPS Deep Learning and Representation Learning Workshop*, 2015.

- [10] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [11] C. Zhang, S. Bengio, and Y. Singer, "Are all layers created equal?" *arXiv preprint arXiv:1902.01996*, 2019.
- [12] C. Chiu, T. N. Sainath, Y. Wu, R. Prabhavalkar *et al.*, "State-of-the-art speech recognition with sequence-to-sequence models," in *Proc. ICASSP*, 2018.
- [13] S. Kim, M. L. Seltzer, J. Li, and R. Zhao, "Improved training for online end-to-end speech recognition systems," in *Proc. INTERSPEECH*, 11 2018. [Online]. Available: <https://arxiv.org/abs/1711.02212>
- [14] L. Xie, J. Wang, Z. Wei, M. Wang, and Q. Tian, "Disturblabel: Regularizing cnn on the loss layer," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [15] L. Lu, M. Guo, and S. Renals, "Knowledge distillation for small-footprint highway networks," in *Proc. ICASSP*, 2017.
- [16] D. Yu, A. Eversole, M. Seltzer, K. Yao, Z. Huang, B. Guenter, O. Kuchaiev, Y. Zhang, F. Seide, H. Wang *et al.*, "An introduction to computational networks and the computational network toolkit," Microsoft Research, Tech. Rep., 2014.