

# Investigation and analysis of hyper and hypo neuron pruning to selectively update neurons during unsupervised adaptation <sup>☆</sup>

Vikramjit Mitra <sup>a,\*</sup>, Horacio Franco <sup>b</sup>

<sup>a</sup> Dept. of Electrical and Computer Eng., University of Maryland, College Park, MD, USA

<sup>b</sup> Speech Technology and Research Laboratory, SRI International, Menlo Park, CA, USA

## ARTICLE INFO

### Article history:

Available online 15 January 2020

### Keywords:

Neural net pruning  
Unsupervised adaptation  
Convolutional neural network speech  
recognition robustness machine learning

## ABSTRACT

Unseen or out-of-domain data can seriously degrade the performance of a neural network model, indicating the model's failure to generalize to unseen data. Neural net pruning can not only help to reduce a model's size but can improve the model's generalization capacity as well. Pruning approaches look for low-salient neurons that are less contributive to a model's decision and hence can be removed from the model. This work investigates if pruning approaches are successful in detecting neurons that are either high-salient (mostly active or hyper) or low-salient (barely active or hypo), and whether removal of such neurons can help to improve the model's generalization capacity. Traditional blind adaptation techniques update either the whole or a subset of layers, but have never explored selectively updating individual neurons across one or more layers. Focusing on the fully connected layers of a convolutional neural network (CNN), this work shows that it may be possible to selectively adapt certain neurons (consisting of the hyper and the hypo neurons) first, followed by a full-network fine tuning. Using the task of automatic speech recognition, this work demonstrates how the removal of hyper and hypo neurons from a model can improve the model's performance on out-of-domain speech data and how selective neuron adaptation can ensure improved performance when compared to traditional blind model adaptation.

© 2020 Elsevier Inc. All rights reserved.

## 1. Introduction

Deep Neural Networks (DNNs) have become the ubiquitous machine learning technique, demonstrating significant performance gains over its peers in almost all applications where they have been tested [1–3], including automatic speech recognition (ASR) [4,5]. DNNs are both data hungry and data sensitive [6].

Neural nets are usually overparameterized with significant redundancy in their number of neurons and the transforms that they learn, given training data [7]. A consequence of information redundancy in neural nets is data over-fitting, where a model ends up learning fine grained structures present within the training data that may not be relevant for the classification task and may not be present in unseen data. Such over-fitting results in models failing to generalize well to unseen data conditions, as a consequence sig-

nificant performance degradation is observed when such models are deployed to unknown datasets as compared to seen or in-domain datasets. To prevent models from over-fitting, well known techniques such as regularization, dropout [8], pruning [9], data augmentation [10] etc., are deployed. Pruning strategies try to minimize data over-fitting by removing less salient neurons from a model while retaining the model's performance on a given dataset. The goal of network pruning is to find neurons in a model whose removal will not degrade the model's performance, while improving its generalization capacity.

Using traditional pruning algorithms, this work investigates the role of neurons that are deemed as less-salient versus the ones that are high-salient. The work investigates the following:

- (1) *Is the traditional approach of pruning out "less salient" neurons the best strategy for performing neural network pruning?*
- (2) *What are the effects of network pruning in each individual layer of a fully connected network?*
- (3) *Is there any benefit in adapting only the "to-be-pruned" neurons at first during unsupervised model adaptation, followed by adapting the whole network?*

The general observation from this study is that there are three main groups of neurons: (1) *less-salient*, or almost inactive or

<sup>☆</sup> This work was performed when the first author was at SRI International. He is currently affiliated with Apple.

\* Corresponding author.

E-mail addresses: [vmitra@umd.edu](mailto:vmitra@umd.edu) (V. Mitra), [horacio.franco@sri.com](mailto:horacio.franco@sri.com) (H. Franco).

'*hypo*' active neurons; (2) *high-salient*, or, very active or, '*hyper*' active neurons, and (3) *medium-salient* neurons, or relevant neurons, whose removal results in catastrophic degradation in a network's performance.

Observations in this paper are based on a speech recognition task, with both '*in-domain*' or '*seen*' data and '*out-of-domain*' or '*unseen*' data as evaluation tasks. The in-domain evaluation data consists of speech recorded in close microphone conditions with varying microphone channels and background degradations. The out-of-domain evaluation data consists of speech recorded using distant microphones, recorded in varying room conditions with different reverberation times and noisy background conditions.

We investigate a new mutual-information based neural network pruning approach, capable of detecting salient neurons leveraging temporal properties of a speech signal and demonstrate its effectiveness in a speech recognition task. Please note that proposing a new pruning approach is not the aim of this paper; instead, understanding which neurons can be pruned and whether there is any benefit in selectively updating such neurons during unsupervised model adaptation is the main goal of this work.

We derive our conclusions based on a state-of-the-art convolutional neural network (CNN) acoustic model [11], focusing on network pruning only for the fully connected layers. Unlike prior work, we investigate layer-wise pruning where we evaluate the impact of pruning on each layer individually.

We investigate if the neurons that are supposed to be pruned, can be used to selectively adapt the neural network model to out-of-domain data. This investigation stems from the fact that certain neurons can be pruned as they have been deemed as "not useful", while the others are retained as they are useful. Rationally, it makes sense to retrain the 'not useful' neurons, while retaining the 'useful' neurons as-is. Our observations show that such selective model adaptation can result in well regularized models that can adapt to out-of-domain data while retaining their performance on in-domain data.

## 2. Data

The acoustic models in this work were trained using the multi-conditioned, noise and channel degraded training data from the 16 kHz Aurora-4 [12] noisy WSJ0 corpus. Aurora-4 data contains six additive noise types with channel-matched and mismatched conditions. It was created from the standard 5K WSJ0 database, containing 15 hours of training data and 0.6 hours (330 utterances) of testing data. The test data is replicated into 14 test sets (0.6 hours each) consisting of two different channel conditions and six different added noises (car; babble; restaurant; street; airport; and train station) in addition to the clean condition. The signal-to-noise ratio (SNR) for the test sets varied between 0 and 15 dB. The evaluation set consists of 5K words. The Aurora-4 test set is used as the in-domain evaluation set in our experiments.

In this work, we treated reverberation as the out-of-domain data condition (not included during training). For adaptation, optimization, and evaluation purposes, we have used the single-microphone subset of the training, dev, and the eval sets distributed with the REVERB 2014 (denoted as REVERB14 in this paper) [13] challenge, respectively. The adaptation set consists of the clean WSJCAM0 [14] data, which was convolved with different room impulse responses (reverberation times from 0.1 to 0.8 sec) and then corrupted with background noise. The performance was evaluated on the dev and test subsets, which contain both real and simulated reverberation conditions. The real data is borrowed from the MC-WSJ-AV corpus [15], which consists of utterances recorded in a noisy and reverberant room. The simulated and the real eval sets contained 1088 and 372 utterances, respectively,

split equally between far and near microphone conditions. Note that none of our experiments used any speaker-level information.

## 3. Acoustic model

Earlier [16] we found that CNN models perform better than the DNNs for Aurora-4 ASR task, hence CNN acoustic models were used in this work. The training alignments consisted of 3125 context-dependent (CD) states, which were generated from a baseline DNN acoustic model. The input acoustic features were formed by using a context of 17 frames (8 frames on either side of the current frame). We observed that time-frequency convolution (using TFCNN [11,17–21]) performed better than 1-D frequency convolution, and hence we have focused on the TFCNN acoustic models for our experiments presented in this paper. The TFCNN architecture is same as in [11,22], where two parallel convolutional layers are used at the input, one performing convolution across time, and the other across frequency on the input filterbank features. The TFCNNs had 75 filters to perform time convolution and 200 filters to perform frequency convolution. For time and frequency convolution, eight bands were used, followed by a max-pooling over three samples after frequency convolution, and five samples for time convolution. The feature maps after both the convolution operations were concatenated and then fed to a fully connected DNN having 2048 nodes and four hidden layers. The TFCNN network was trained using four initial iterations with a constant learning rate of 0.008, followed by learning rate halving based on cross-validation error. Training stopped when no further reduction in cross-validation error was noted or when it started to increase. Backpropagation was performed by using stochastic gradient descent with a mini-batch size of 256 examples. Note that each model has been trained three times to interpret the uncertainty in the model performance. In our experiments we provide the mean WER from the multiple iterations and their standard deviations are reported within parenthesis in tables and as error bars in figures. Note that the baseline acoustic model was trained from a random initial seed, and that baseline model has been used as the seed model for all our experiments.

We used gammatone filter-bank energies (GFBs) as the acoustic feature for our experiments. GFBs were generated using a bank of 40 time-domain gammatone filters, using an analysis window of 26 millisecond and a frame rate of 10 millisecond. The gammatone subband powers were compressed using 15th root.

## 4. Neural network pruning

One main goal of this work is to understand how layer-wise pruning affects performance on out-of-domain data and validate that across multiple pruning strategies. We explored three pruning strategies and they are briefly described below:

(1) *Magnitude based pruning* (MBP): assumes that small weights are less important [23–25]. Two factors are typically used for determining redundant neurons: consuming energy and weight power [20]. In this work, we ascertained a neuron's saliency by using its weight's power [24].

(2) *Optimal brain surgeon* (OBS) is a sensitivity-based approach that attempts to find the contribution of each weight or node in the network and then prunes the weight or node that have the least effect on the objective function. One of the celebrated sensitivity-based pruning algorithms is Optimal Brain Damage (OBD) [26], that approximates the measure of 'saliency' of a weight by estimating the second derivative of the network output error with respect to that weight. OBS [27] relaxed some of the assumptions in OBD. In this work we have used OBS to ascertain neural saliency of each layer.

(3) *Mutual independence* (MI) based pruning leverage relationships between input and output of a neuron to predict the neuron's saliency. One such approach is mutual information based [28,29], which uses singular value decomposition to analyze the hidden unit activation covariance matrix, where the rank of the covariance matrix determines the optimal number of hidden units.

In this work we investigated a cross-correlation based MI measure to estimate the input-output relationship of a neuron to determine the saliency of that neuron, given an input feature set. The cross-correlation analysis can be performed using a subset of the training data after neural network training.

Let  $a_{n,l}$  be the  $n$ th neuron of the  $l$ th layer, whose inputs are:

$$[x_{1,l-1}[t], x_{2,l-1}[t], \dots, x_{p,l-1}[t], \dots, x_{N,l-1}[t]]$$

where  $x_{p,l-1}[t]$  is the activation from the  $p$ th neuron of the preceding  $l-1$ th layer at time instant  $t$ . Let the activation from the  $a_{n,l}$  at  $t$  be  $x_{n,l}[t]$ , assuming that  $l-1$ th layer had  $N$  neurons, the mean absolute cumulative cross-correlation between the input and the output activations for the neuron  $a_{n,l}$  over a time window  $q$  at time  $t$  will be:

$$\hat{r}_{n,l}[t] = 1/N \sum_{p=1}^N \left| \sum_{k=-(q/2-1)}^{q/2} x_{n,l}[t] x_{p,l-1}[t_k] \right| \quad (1)$$

Saliency of neuron  $a_{n,l}$  of layer  $l$  is determined by the value of  $\hat{r}_{n,l}[t]$ ; higher the value the more salient the neuron is and vice versa. Note that, this measure is particularly suitable for quasi-periodic pseudo-stationary signals like speech, that retain a certain degree of temporal correlation. The measure in (1) indicates that if the inputs to the neuron  $n$  is not correlated to the output from the neuron  $n$ , then neuron  $n$  is less salient, in the sense that is in not conveying information to the succeeding layers.  $\hat{r}_{n,l}[t]$  is used to estimate saliency in the MI based pruning investigated in this paper.

Note that while pruning, we rank-sorted the neurons based on their saliency, where the top salient neurons are selected as high-salient or hyper-active neurons and the low-salient neurons are selected as low-salient or hypo-active neurons.

## 5. Neural network pruning

We investigated layer-wise pruning, focusing only on the fully connected hidden layers, by pruning 2% to 12% in steps of 2%. We investigate layer-wise pruning in three possible ways:

- (a) pruning neurons with low saliency,
- (b) pruning neurons with high saliency,
- (c) pruning neurons with medium saliency (i.e., the neurons that are in between (a) and (b)).

Note that during layer-wise pruning a 2% pruning means pruning 2% of the neurons in that layer, for example if the number of neurons in that layer is 2048, a 2% pruning of that layer would indicate removing 41 neurons from that layer.

### 5.1. Pruning of low salient (hypo) neurons

These are the neurons that are least contributive to information flow within the neural net and have been the candidates for pruning in traditional approaches. We investigated MBP, OBS and MI based pruning of neurons in each fully connected layers of the baseline TFCNN model and analyzed their role on both in-domain and out-of-domain data. The word error rates (WERs) on the Aurora-4 evaluation set, real and simulated dev sets of REVERB14 data is given in Table 1, when 2% pruning is performed only on the 1st layer of the fully connected network. Fig. 1 shows the outcome of pruning on the 1st hidden layer, demonstrating

**Table 1**

WERs (in %) from the baseline (not pruned) model and the 2% pruned (fully connected layer 1 only) acoustic model when evaluated on the Aurora-4 and REVERB14 dev (simulated and real condition) set.

Pruning	Layer	Pruning %	Aurora-4	REVERB14	
				Simulated	Real
None	–	0	9.1	39.3	42.4
MBP	1	2	9.2 (0.0)	38.5 (0.3)	41.0 (0.4)
OBS	1	2	9.2 (0.1)	38.2 (0.1)	40.1 (0.7)
MI	1	2	9.2 (0.1)	36.6 (0.3)	38.9 (0.2)

that it did not significantly impact the model's performance on in-domain data, with up to 6% relative deterioration at 10% pruning level; where MBP showed the least deterioration out of the three pruning approaches investigated.

Fig. 2 shows the effect of 1st hidden layer pruning on out-of-domain data. Interestingly, Fig. 2 shows that pruning resulted in performance improvement for out-of-domain data (compared to the unpruned baseline model), with maximum relative reduction of 9.6% in WER obtained from MI at 8% pruning level.

Note that Figs. 1 and 2 correspond to pruning only the 1st fully-connected hidden layer. Figs. 3 and 4 show the effect on the performance after pruning the low-salient neurons from the 2nd and the 3rd fully-connected hidden layer. The performance on in-domain data was found to be similar (as shown in Fig. 1) for the first two fully connected layers, for up to 12% pruning. The 3rd hidden layer (the prefinal layer) was found to be quite sensitive to pruning (as observed from Fig. 4), where pruning beyond 6% neurons resulted in more than 10% relative increased in WER for in-domain data.


Several interesting observations can be made from Figs. 2, 3 and 4. While MI always performed the best from pruning the 1st and 3rd hidden layers, OBS performed the best for pruning the 2nd hidden layer. Note that the 3rd layer was the pre-final layer, hence pruning beyond 8% was found to adversely affect the output layer decisions.

Note that MI uses temporal correlations, and as speech is a quasi-stationary signal, the acoustic features across time have temporal dependencies. The output targets (in this case the senone labels) also have temporal dependencies. The above fact may justify why the performance from MI-based pruning was better across the different layers of the network, compared to the other two approaches. From the above figures the following observations can be made:


- (I) Pruning on first and second hidden layer tend to increase the generalization capacity of the network, which may indicate that there is more redundancy in the hidden layers than in the pre-final layer.
- (II) The pre-final layer seems to be quite sensitive to pruning, where pruning beyond 6% may significantly reduce the performance of the model. Interestingly, for this case MI-based pruning did not show any performance degradation.
- (III) Removal of low-salient (hypo-neurons) from the network results in improved performance on out-of-domain data-set, while maintaining the performance on in-domain data. Indicating that the model's generalization capacity can be improved by pruning of the hypo neurons.

### 5.2. Pruning of hi-salient (hyper) neurons

We investigated the role of the high salient (or hyper active) neurons, by layer-wise pruning of such neurons and observing model performance for both in-domain and out-of-domain data. We explored all the three pruning techniques in a layer-wise fashion like the steps in 5.2.



**Fig. 1.** WER versus low-salient (hypo-) neuron pruning levels in fully-connected layer-1 for in-domain data (Aurora-4). Red dotted line shows the same for the un-pruned baseline model. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)



**Fig. 2.** WER versus low-salient (hypo) neuron pruning levels in fully-connected layer-1 for out-of-domain data (REVERB14-real dev set). Red dotted line shows the same for the un-pruned baseline model.


The WERs on the Aurora-4 evaluation set, and the real dev set of REVERB14 data are shown in Figs. 5 to 8, when pruning is performed from 2% to 12% in steps of 2% for the 1st, 2nd and the 3rd fully connected hidden layers.

Figs. 5 and 6 show that MI is the most sensitive amongst the three for hyper-neuron pruning. The model performance reduces for MI after 4% pruning of hyper-neuron. Whereas for OBS and MBP, up to 10% pruning of hyper-neuron was possible, which gave better performance on the out-of-domain data, where WER reduction as high as 6.8% relative to the baseline was observed.


Fig. 7 shows that for 2nd hidden layer, all three pruning techniques resulted in WER reductions, however Fig. 8 shows that the same pruning strategies on the pre-final layer resulted in significant deterioration of performance beyond a pruning threshold of 4% hyper neurons.

From Figs. 5 to 8 the following observations can be made:

(I) Pruning of high-salient or hyper neurons can retain baseline performance for in-domain data and can improve performance for out-of-domain data. Hyper neurons are the ones that are more likely to fire given an input feature, hence they may not be contributing to the discriminative power of the network, and thus



**Fig. 3.** WER versus low-salient (hypo) neuron pruning levels in fully-connected layer-2 for out-of-domain data (REVERB14-real dev set). Red dotted line shows the same for the un-pruned baseline model.



**Fig. 4.** WER versus low-salient (hypo) neuron pruning levels in fully-connected layer-3 for out-of-domain data (REVERB14-real dev set). Red dotted line shows the same for the un-pruned baseline model.

can be removed without a significant loss in performance. We observed that the removal of up to 4% of the top hyper neurons can retain in-domain performance while improving the generalization capacity of the network to out-of-domain data. To the best of our knowledge, this is the first claim where hyper active neuron pruning was explored and was found to be useful.

(II) Salient neurons defined by MI based approach are found to be more essential to the neural net's performance, than from the other two techniques. Removal of MI-based salient neurons was found to result in significant deterioration of the network's per-

formance on in-domain data, as compared to OBS and MBP, as observed from Fig. 5.

(III) Removal of salient-neurons from the pre-final layer was found to have deleterious effect, where pruning beyond 4% of top salient features, resulted in performance degradation across all three pruning methods for both in-domain and out-of-domain data.

### 5.3. Pruning medium-salient neurons

Finally, we explored pruning of the medium salient neurons that are neither the lowest 12% salient neurons nor the top 12%




Fig. 5. WER versus hyper-neuron pruning levels in fully-connected layer-1 for in-domain data (Aurora-4). Red dotted line shows the same for the un-pruned baseline model.




Fig. 6. WER versus hyper-neuron pruning levels in fully-connected layer-1 for out-of-domain data (REVERB14-real dev set). Red dotted line shows the same for the un-pruned baseline model.


salient neurons. Exploring layer-wise pruning of mid-salient neurons resulted in catastrophic performance failure (with WER reaching as high as 100%) for both in-domain and out-of-domain data, indicating that mid-salient neurons are crucial for a neural net's performance.

While pruning the mid-salient neurons we investigated pruning of the neurons in steps of 10%, from 5% pruning to pruning of all mid-salient neurons; and we observed the baseline WER to range from 16.5% to 100%. The results indicated that with 5% pruning of mid-salient neurons the baseline performance deteriorates by more than 80%.


#### 5.4. Pruning of both hyper- and hypo- neurons

Given the observations in sections 5.1, 5.2 and 5.3, we investigated layer-wise pruning of both hyper- and hypo- neurons and pruning across multiple layers. Based on findings from sections 5.1 and 5.2, we investigated 8% pruning of the hypo-neurons and 4% pruning of the hyper-neurons on layer 1 (that resulted in an overall pruning of 4% across all the three fully connected layers), and the results are shown in Table 2. Table 2, shows that pruning both the hyper and hypo neurons in fully connected layer 1, resulted in retaining in-domain performance, while improving the performance on out-of-domain data-set, indicating that such pruning strategy





**Fig. 7.** WER versus hyper-neuron pruning levels in fully-connected layer-2 for out-of-domain data (REVERB14-real dev set). Red dotted line shows the same for the un-pruned baseline model.



**Fig. 8.** WER versus hyper-neuron pruning levels in fully-connected layer-3 for out-of-domain data (REVERB14-real dev set). Red dotted line shows the same for the un-pruned baseline model.

helped to improve the model's generalization capacity to unseen data conditions. Table 2 shows that a 4% pruning of hyper and hypo neurons (based on MI) resulted in WER reduction relative to the baseline of 8.7% and 7.5% for simulated and real Reverb14 dev data sets respectively.

Next, we investigated the pruning of hyper- and hypo- neurons across multiple layers. We investigated (a) pruning the three fully connected hidden layers and (b) pruning the first two fully connected layers only (ignoring the prefinal layer, as pruning of that layer was not found to be useful from Figs. 4 and 8). Table 3 shows the results from pruning across multiple layers. Note that for the

1st two fully connected layers 8% and 4% hypo and hyper neurons were pruned respectively, and for the pre-final layer 2% hyper and 2% hypo neurons were pruned respectively. Table 3 shows that pruning across all the three layers had more deleterious effect than pruning neurons only from the first two hidden layers, indicating that existing blind pruning techniques may be missing out on potential benefits of selective layer-wise pruning. Note that overall 8% pruning was performed for the experimental results shown in Table 3 compared to 4% overall pruning for the results shown in Table 2, which justifies the slightly worse out-of-domain performance in Table 3.

**Table 2**

WERs from the baseline (not pruned) model and the 8% hypo-neuron and 4% hyper-neuron pruned (deep layer 1 only  $\approx$  4% pruning across 3 fully-connected layers) acoustic model when evaluated on the Aurora-4 and REVERB14 dev-set.

Pruning	Layer	Pruning %	Aurora-4	REVERB14	
				Simulated	Real
None	–	0	9.1	39.3	42.4
MBP	1	4	9.2 (0.1)	38.1 (0.3)	41.4 (0.9)
OBS	1	4	9.6 (0.3)	37.5 (0.5)	40.4 (0.8)
MI	1	4	9.3 (0.1)	35.9 (0.2)	39.2 (0.5)

**Table 3**

WERs from the baseline (not pruned) model and multi-layer pruned (pruning percentage computed over the 3 hidden layers) acoustic model when evaluated on the Aurora-4 and REVERB 2014 dev-set. Percent pruning computed for the entire network.

Pruning	Layer	Pruning %	Aurora-4	REVERB14	
				Simulated	Real
None	–	0	9.1	39.3	42.4
MBP	1	8	9.1 (0.2)	39.1 (0.1)	43.1 (0.7)
OBS	1	8	9.7 (0.3)	39.3 (0.7)	42.1 (0.2)
MI	1	8	9.4 (0.1)	36.9 (0.3)	41.2 (0.4)
MBP	1-2	8	9.7 (0.3)	39.6 (0.3)	43.7 (0.8)
OBS	1-2	8	9.9 (0.5)	37.0 (0.5)	41.7 (0.9)
MI	1-2	8	9.7 (0.3)	36.7 (0.3)	41.0 (0.5)
MBP	1-2-3	9	10.4 (0.2)	41.2 (0.4)	44.2 (0.8)
OBS	1-2-3	9	10.5 (0.4)	39.0 (0.5)	44.3 (0.9)
MI	1-2-3	9	9.9 (0.5)	37.8 (0.3)	42.6 (0.6)

**Table 4**

WERs from the baseline (not pruned) model and multi-layer pruned (pruning percentage computed over the 3 hidden layers) acoustic model when evaluated on the REVERB 2014 test-set. Percent pruning computed for the entire network.

Pruning	Layer	Pruning %	REVERB14	
			Simulated	Real
None	–	0	37.8	46.9
MBP	1	8	37.2	47.5
OBS	1	8	36.6	46.8
MI	1	8	35.5	46.7
MBP	1-2	8	36.0	44.9
OBS	1-2	8	35.1	43.8
MI	1-2	8	<b>34.0</b>	<b>43.4</b>
MBP	1-2-3	9	37.3	46.9
OBS	1-2-3	9	35.1	44.6
MI	1-2-3	9	34.4	44.2

Table 3 shows that pruning over first two fully-connected hidden layers gave much better performance compared to pruning across all the fully connected layers or pruning only the first hidden layer. A relative reduction in WER (for the MI based pruning) was found to be 6.7% and 3.3%, for the simulated and real sets of out-of-domain (REVERB14 dev) data-set.

Table 4 presents recognition results from the out-of-domain (REVERB14) data-set, where observations are similar to that in Table 3. Pruning across the first two fully connected layers were found to be more useful than pruning across all three layers or the first hidden layer alone. A relative reduction in WER (for the MI based pruning) was found to be 10% and 7.5%, for the simulated and real sets of out-of-domain (REVERB14 test) data-set.

Following the observations from Tables 3 and 4, we can state the following:

- (a) Pruning both the hyper and hypo neurons can help to improve the generalization capacity of the network.
- (b) The best results on out-of-domain data-set was observed from pruning the fully connected hidden layers without the pre-final layer. The pre-final layer was found to be quite sensitive to pruning for the given task of speech recognition.

**Table 5**

WERs from the baseline (not pruned) model, adapted baseline models: Models A, B, C and D (as outlined above), when evaluated on the Aurora-4 and REVERB14 dev set.

Models	Layer update	Aurora-4	REVERB14	
			Simulated	Real
Baseline	0	9.1	39.3	42.4
Model-A	All	15.3	24.8	33.2
Model-B	1-2	12.6	23.9	32.8
Model-C	1-2	9.5	22.2	32.2
Model-D	All	9.5	22.0	31.5

Network pruning relies on the fact that a subset of neurons may not have learned reliable information, or, may have learned redundant information, hence preventing the model to generalize well to unseen data conditions. As a consequence, removing such redundant neurons from the network may not cause a significant loss in performance. Table 3 shows that an 8% pruning of hyper and hypo neurons across the first two fully connected hidden layers resulted in a relative WER increase of 6.5% compared to the unpruned baseline model.

Instead of pruning the redundant neurons, they can be updated during unsupervised model adaptation. Typically, during model adaptation, all neurons in one specific layer, or a subset of layers or the whole model are blindly updated, given an adaptation set; where over-fitting is prevented using a suitable regularization parameter. In this work we investigate if there is any benefit in performing initial adaptation of only the redundant neurons (that are supposed to be pruned) and then fine tuning the whole model. Such an unsupervised adaptation technique will focus on learning the new patterns from out-of-domain data through updating the redundant neurons only.

## 6. Acoustic model adaptation

For acoustic model adaptation, we have used the REVERB14 training data-set as the unsupervised adaptation data. The baseline TFCNN acoustic model was used to generate labels for the unsupervised adaptation set and those labels were used to retrain the network. During adaptation, model parameters were updated with an L2 regularization of 0.001 and an initial learning rate of 0.004, with the learning rate halving at every step and a maximum epoch of 10.

Table 5 shows the results from the:

- (a) **Baseline model:** Unpruned TFCNN acoustic model.
- (b) **Model A:** Baseline model after traditional unsupervised adaptation (blind model update), where all parameters of the network were updated.
- (c) **Model B:** Baseline model after selective neuron-update based unsupervised model adaptation, where only the redundant neurons (neurons that can be pruned) in the network were updated.
- (d) **Model C:** obtained by fine-tuning of Model B, where the adaptation set is augmented with part of the original Aurora-4 training set ( $\approx$  50%) and all of the model parameters were updated.
- (e) **Model D:** obtained by updating all parameters of the baseline model using a mix of the REVERB14 training set and the original Aurora-4 training set ( $\approx$  50%).

Evaluation of the above five models were performed on the Aurora-4 test set, REVERB-14 real dev and test sets. For unsupervised adaptation, we have used data-selection using the procedure outlined in [30]. Table 5 shows that the selective neuron-update based unsupervised adaptation (Model B) gave better performance than traditional (blind model update) unsupervised adaptation (Model A). Both cases of selective and blind model update resulted in reduction in WER for the out-of-domain data (simulated and real REVERB-2014 dev sets), however, that came at a cost of



**Table 6**

WERs from the baseline (not pruned) model, adapted baseline models: Models A, B, C and D (as outlined above), when evaluated on the REVERB 2014 test set.

Models	Layer update	REVERB14	
		Simulated	Real
Baseline	0	37.8	46.9
Model-A	All	22.9	37.1
Model-B	1-2	22.4	35.9
Model-C	1-2	20.8	35.7
Model-D	All	20.6	34.7

increase in WER on in-domain data (Aurora-4), due to the phenomena commonly known as the *catastrophic forgetting*. Table 5 shows that the blind model adaptation had more pronounced performance degradation on in-domain data, which is 68% relative to the baseline, as compared to the selective neuron-update based unsupervised adaptation, where performance degradation was 38% relative to the baseline. Hence, Model B suffered less from *catastrophic forgetting* than model A, which can be attributed to the selective model update. Based on the observation from Tables 3 and 4, the selective unsupervised adaptation was performed only on the hyper and hypo neurons of layers 1 and 2, based on the decisions from the MI pruning technique.

The deterioration in the performance on the in-domain data is expected for the selective neuron-update based model adaptation, as only a part of the network gets updated during unsupervised adaptation. Such a selective parameter update can result in altering some network connections that may have had learned patterns in in-domain training data. Model C overcomes such a shift, through the fine-tuning of the Model B by using a part of in-domain training data in addition to unsupervised adaptation data to update the model parameters. After the fine-tuning, Model C demonstrated in-domain performance almost as-good-as the baseline, which can be observed from the final row of Table 5.

Additionally, Model D reflects the case where all the parameters of the baseline model is updated with a mix of REVERB14 training set and the Aurora-4 training set. Results from Model D in Table 5 shows that the performance of Model C and Model D are very close to each other, where Model D has a mildly lower WER for REVERB14 eval sets than the Model C. Model-D has the advantage of having all parameter update given the adaptation set and the original training set compared to the selective model update followed by fine tuning in Model C. However, the interesting observation is that the fine-tuning of the model after selective adaptation resulted in a model (Model C) which performed almost as-good-as the blind-model update using both in-domain (Aurora-4) and out-of-domain (REVERB14) data-sets. However, both models C and D assumes the existence of the original training data, which model B does not, hence Table 5 clearly shows that in absence of the original training data the selective model update (Model B) can reduce the effect of *catastrophic forgetting* as typically observed in blind model updates (model A). In presence of the original training data both models (selectively trained or not) seems to converge to similar performance.

Table 6 shows the results from the baseline model, models A, B, C and D for the REVERB14 test set. Similar to Table 5, both models A and B show lower WERs for the simulated and real test sets compared to the baseline model. Model B shows better performance than the model A, indicating the effectiveness of the selective-neuron update in unsupervised adaptation. The fine-tuning step resulted in even better performance as observed from the WER obtained from Model C in Table 6. Fine-tuning lowered the WER by 7.1% and 0.5% for the simulated and real test sets respectively, relative to Model B. Selective neuron-update followed by fine-tuning based unsupervised adaptation (model C) resulted in similar performance as the blind model update (Model D) using

both REVER14 and Aurora-4 adaptation sets, bolstering the observation from Table 5 showing that the fine-tuning step helped the model to converge to blind model update, where both in-domain and out-of-domain data-sets were used for adaptation. However, in case of the absence of in-domain data for adaptation, the Model B clearly demonstrates that it is relatively resilient to *catastrophic forgetting* as compared to the blind model update.

## 7. Observations

The goal of this study has been to understand the following:

- How does neural network pruning help to improve the model's generalization capacity?
- What are the effects of network pruning in each individual layer of a fully connected network?
- How does the model's generalization capacity to unseen data vary when we prune the low-salient neurons (that is neurons that are barely active or hypo neurons) as opposed to the high salient neurons (that is neurons that are mostly active or hyper neurons), and if there is any trade-off between choosing one over the other during the network pruning step?
- If there is any benefit of adapting the to-be-pruned neurons during unsupervised adaptation?

Based on the results obtained from our studies, we can make the following observations:

- We observed a significant improvement in an acoustic model's generalization capacity to out-of-domain data, when pruning was performed at different layers of the fully connected network. We found that pruning of up to 6% of the neurons can be performed which resulted in retaining in-domain performance, while improving performance on out-of-domain data. Overall an average of 4.5% relative reduction in WER was observed for out-of-domain data, while no significant change in WER was observed for the in-domain data, from pruning 6% of the low salient neurons from each of the three layers of the fully connected layer using MI based pruning.
- We observed that each individual layer behaved differently to pruning for the given task of acoustic modeling. Pruning of the first two fully connected layer resulted in significant reduction of WER (with over 97% confidence), however pruning the pre-final fully connected layer did not result in such significant reduction in WER. The justification of such an observation may be the fact that more redundancy may prevail in the middle layers than the pre-final layer, hence pruning the latter had more deleterious effect on the overall performance of the model.
- We observed that pruning of low-salient neurons was quite effective in improving the model's generalization capacity to out-of-domain dataset without sacrificing significant performance degradation on in-domain data. Additionally, we observed that hyper-neurons or overly-active neurons can also be pruned, with similar results.

Traditionally, the literature has focused on only the pruning of low-salient neurons and this work shows that hyper active neurons can also be pruned from the model without significantly impairing the model's performance. Experimental results from the pruned acoustic models in our studies demonstrated substantial improvement in recognition results on out-of-domain data as compared to the unpruned model.

- Finally, we investigated if the to-be-pruned neurons can be selectively updated during unsupervised model adaptation, instead of pruning them out. Results indicated that updating such neurons during unsupervised adaptation results in well regularized networks, that can learn additional information from the out-of-domain adaptation data without significant drop in performance on the in-domain data. Table 6 shows that the selective neuron update-based model adaptation (model B) resulted in 40.7% and

23.5% relative reduction in out-of-domain WER compared to the baseline for simulated and real test sets of REVERB 2014, respectively. The blind model adaptation (model A) provided a 39.4% and 20.9% relative reduction in out-of-domain WER compared to the baseline for simulated and real test sets of REVERB 2014, respectively. The deterioration in in-domain WER was 38.5% in case of the selective neuron update based unsupervised adaptation (model B) as opposed to 68.1% from the blind model update (model A). The fine-tuning stage (updating all the model parameters) after the selective model update resulted in 22.3% relative reduction in WER for the out-of-domain data with a relative increase in WER of only 4% for the in-domain data. The improvement in ASR performance from the two-stage unsupervised adaptation was statistically significant when compared to the more traditional blind unsupervised adaptation technique.

## 8. Conclusion

In this work, we investigated neural net pruning, to better understand which neurons can be pruned to improve the generalization capacity of a network to out-of-domain data. We observed that pruning can help in generalization and found that pruning at different layers may have different impacts, where aggressive pruning toward the final layers may not be desirable. We observed that the fully connected layers have broadly three groups of neurons: hypo-neurons, hyper-neurons, and the informative neurons (which are in between the hyper and hypo neurons w.r.t their saliency). Literature primarily focused on pruning off the hypo neurons, but we observed that pruning both hypo and hyper neurons can not only help in reducing the model size, but also improve the model's generalization capacity.

Finally, we investigated whether instead of pruning the hyper and hypo neurons, we can selectively update those neurons during unsupervised model adaptation and found such strategy to be more successful compared to blind model update. We observed that selective model update gave better performance on out-of-domain data and diverged less from the in-domain data, compared to the blind model update approach. Finally, a fine-tuning adaptation step using a combination of in-domain and out-of-domain data was found to provide in-domain performance almost as-good-as the in-domain baseline model and a significant relative reduction in out-of-domain WER of more than 22% (for the real test set of REVERB 2014) compared to the baseline model. The findings from this study indicates that during unsupervised adaptation, it may be beneficial to first update the neurons that are deemed as prunable by standard pruning techniques and then perform fine-tuning by updating all the model parameters.

In this work we have focused only on pruning the fully-connected layers, leaving the convolutional layers untouched. From our preliminary investigation we did not find pruning the convolutional layers as-promising-as pruning the fully-connected layer. In future studies, we plan to investigate more on pruning of the convolutional layers, to see if we can find more room for improvement after the pruning of the fully connected layers.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgment

The authors would like to acknowledge Dimitra Vergyri, the reviewers and the Editor-in-chief of DSP for their helpful thoughts, comments and suggestions.

## References

- [1] A. Krizhevsky, I. Sutskever, G.E. Hinton, Imagenet classification with deep convolutional neural networks, in: *Advances in Neural Information Processing Systems*, 2012, pp. 1097–1105.
- [2] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [3] T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient estimation of word representations in vector space, *arXiv preprint arXiv:1301.3781*.
- [4] A.-r. Mohamed, G.E. Dahl, G. Hinton, Acoustic modeling using deep belief networks, *IEEE Trans. Audio Speech Lang. Process.* 20 (1) (2011) 14–22.
- [5] F. Seide, G. Li, D. Yu, Conversational speech transcription using context-dependent deep neural networks, in: *Twelfth Annual Conference of the International Speech Communication Association*, 2011.
- [6] F. Grézil, E. Egorova, M. Karafiát, Further investigation into multilingual training and adaptation of stacked bottle-neck neural network structure, in: *2014 IEEE Spoken Language Technology Workshop (SLT)*, IEEE, 2014, pp. 48–53.
- [7] C. Louizos, M. Welling, D.P. Kingma, Learning sparse neural networks through L0 regularization, *arXiv preprint arXiv:1712.01312*.
- [8] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: a simple way to prevent neural networks from overfitting, *J. Mach. Learn. Res.* 15 (1) (2014) 1929–1958.
- [9] J. Hertz, A. Krogh, R.G. Palmer, H. Horner, Introduction to the theory of neural computation, *Phys. Today* 44 (1991) 70.
- [10] M. Karafiát, F. Grézil, L. Burget, I. Szöke, J. Černocký, Three ways to adapt a cts recognizer to unseen reverberated speech in but system for the aspire challenge, in: *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [11] V. Mitra, H. Franco, Time-frequency convolutional networks for robust speech recognition, in: *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, IEEE, 2015, pp. 317–323.
- [12] G. Hirsch, Experimental framework for the performance evaluation of speech recognition front-ends on a large vocabulary task, *ETSI STQ Aurora DSR Working Group*, Dec. 2002.
- [13] K. Kinoshita, M. Delcroix, T. Yoshioka, T. Nakatani, E. Habets, R. Haeb-Umbach, V. Leutnant, A. Sehr, W. Kellermann, R. Maas, et al., The reverb challenge: a common evaluation framework for dereverberation and recognition of reverberant speech, in: *2013 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, IEEE, 2013, pp. 1–4.
- [14] T. Robinson, J. Fransen, D. Pye, J. Foote, S. Renals, WSJCAMO: a British English speech corpus for large vocabulary continuous speech recognition, *1995 International Conference on Acoustics, Speech, and Signal Processing*, vol. 1, IEEE, 1995, pp. 81–84.
- [15] M. Lincoln, I. McCowan, J. Vepa, H.K. Maganti, The multi-channel wall street journal audio visual corpus (mc-wsj-av): specification and initial experiments, in: *IEEE Workshop on Automatic Speech Recognition and Understanding*, 2005, IEEE, 2005, pp. 357–362.
- [16] V. Mitra, W. Wang, H. Franco, Y. Lei, C. Bartels, M. Graciarena, Evaluating robust features on deep neural networks for speech recognition in noisy and channel mismatched conditions, in: *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- [17] E. Yilmaz, V. Mitra, C. Bartels, H. Franco, Articulatory features for ASR of pathological speech, *arXiv preprint arXiv:1807.10948*.
- [18] E. Yilmaz, V. Mitra, G. Sivaraman, H. Franco, Articulatory and bottleneck features for speaker-independent ASR of dysarthric speech, *Comput. Speech Lang.* 58 (2019) 319–334.
- [19] V. Mitra, H. Franco, C. Bartels, J. van Hout, M. Graciarena, D. Vergyri, Speech recognition in unseen and noisy channel conditions, in: *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2017, pp. 5215–5219.
- [20] V. Mitra, G. Sivaraman, H. Nam, C. Espy-Wilson, E. Saltzman, M. Tiede, Hybrid convolutional neural networks for articulatory and acoustic information based speech recognition, *Speech Commun.* 89 (2017) 103–112.
- [21] V. Mitra, H. Franco, R.M. Stern, J. Van Hout, L. Ferrer, M. Graciarena, W. Wang, D. Vergyri, A. Alwan, J.H. Hansen, Robust features in deep-learning-based speech recognition, in: *New Era for Robust Speech Recognition*, Springer, 2017, pp. 187–217.
- [22] V. Mitra, J. van Hout, W. Wang, C. Bartels, H. Franco, D. Vergyri, A. Alwan, A. Janin, J.H. Hansen, R.M. Stern, et al., Fusion strategies for robust speech recognition and keyword spotting for channel-and noise-degraded speech, in: *INTERSPEECH*, 2016, pp. 3683–3687.
- [23] W. Wan, S. Mabu, K. Shimada, K. Hirasawa, J. Hu, Enhancing the generalization ability of neural networks through controlling the hidden layers, *Appl. Soft Comput.* 9 (1) (2009) 404–414.
- [24] M. Hagiwara, A simple and effective method for removal of hidden units and weights, *Neurocomputing* 6 (2) (1994) 207–218.
- [25] J. Sietsma, Neural net pruning-why and how, in: *Proceedings of International Conference on Neural Networks*, San Diego, CA, 1988, vol. 1, 1988, pp. 325–333.
- [26] Y. LeCun, J.S. Denker, S.A. Solla, Optimal brain damage, in: *Advances in Neural Information Processing Systems*, 1990, pp. 598–605.

- [27] B. Hassibi, D.G. Stork, G.J. Wolff, Optimal brain surgeon and general network pruning, in: IEEE International Conference on Neural Networks, IEEE, 1993, pp. 293–299.
- [28] H.-J. Xing, B.-G. Hu, Two-phase construction of multilayer perceptrons using information theory, IEEE Trans. Neural Netw. 20 (4) (2009) 715–721.
- [29] Z. Zhang, J. Qiao, A node pruning algorithm for feedforward neural network based on neural complexity, in: 2010 International Conference on Intelligent Control and Information Processing, IEEE, 2010, pp. 406–410.
- [30] V. Mitra, H. Franco, Interpreting dnn output layer activations: a strategy to cope with unseen data in speech recognition, in: 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, 2018, pp. 5724–5728.

**Vikramjit Mitra** is a Senior Member of the IEEE, Senior Research Scientist at Apple and a Research Affiliate at University of Maryland - College Park, who is working on speech science and machine learning for human-machine interactions. Previously he worked as an Advanced Research Scientist at SRI International's Speech Technology and Research Laboratory from 2011 to 2017. He received his PhD in Electrical Engineering from University of Maryland, College Park in 2010. His research interests include speech for health applications, robust sig-

nal processing for noise/channel/reverberation, speech recognition, production/perception-motivated signal processing, information retrieval, machine learning, and speech analytics. He has served as the PI/co-PI of several projects funded by NSF and has worked on research efforts funded by DARPA, IARPA, AFRL, NSF and Sandia National Laboratories. He is a senior member of the IEEE and an affiliate member of the Speech and Language processing technical committee (SLTC), and he has served on the scientific committees of several workshops and technical conferences.

**Horacio Franco** is Chief Scientist at the Speech Technology and Research Laboratory at SRI International. He received his Doctor in Engineering degree in 1996, from the University of Buenos Aires, Argentina. Since joining SRI in 1990, he has contributed to the development of SRI's speech technology in several aspects of acoustic modeling and recognizer architectures, speech technology for language learning, speech translation systems, and noise robustness for speech recognition. Dr. Franco has been the principal investigator for several DARPA funded projects, some of which led to successful technology transitions.