AUTOMATIC NODE SELECTION FOR DEEP NEURAL NETWORKS USING GROUP LASSO REGULARIZATION

Tsubasa Ochiai^{1,*}, Shigeki Matsuda¹, Hideyuki Watanabe², and Shigeru Katagiri¹

¹Graduate School of Science and Engineering, Doshisha University, Kyoto, Japan ²Advanced Telecommunications Research Institute International, Kyoto, Japan * eup1105@mail4.doshisha.ac.jp

ABSTRACT

We examine the effect of the Group Lasso (gLasso) regularizer in selecting the salient nodes of Deep Neural Network (DNN) hidden layers by applying a DNN-HMM hybrid speech recognizer to TED Talks speech data. We test two types of gLasso regularization, one for outgoing weight vectors and another for incoming weight vectors, as well as two sizes of DNNs: 2048 hidden layer nodes and 4096 nodes. Furthermore, we compare gLasso and L2 regularizers. Our experiment results demonstrate that our DNN training, in which the gLasso regularizer was embedded, successfully selected the hidden layer nodes that are necessary and sufficient for achieving high classification power.

Index Terms— Deep neural networks, Group Lasso regularization, Speech recognition

1. INTRODUCTION

Motivated by the recent rise of Deep Neural Networks (DNNs) for speech recognition [1], we studied the utility of DNN-based speaker adaptation [2, 3]. A DNN's high capability is mainly derived from a deep layer structure and a massive amount of network weight parameters. Therefore, roughly speaking, the larger and deeper the network, the higher its performances are. However, a large network is obviously undesirable from the viewpoints of computational load and memory size; it is also unfavorable from the viewpoint of controlling training robustness to unseen data. To meet this requirement for finding a small, necessary, and sufficient DNN structure, several approaches have reshaped the network structure [4, 5, 6] or pruned the network nodes [7]. However, these methods assumed retraining or adapting a size-reduced network for high discriminative power. Therefore, a need obviously exists for developing training methods that can automatically (without additional retraining and adaptation) find a small but sufficient DNN structure.

Seeking such development, we propose in this paper a new training scheme that embeds group Lasso (gLasso) regularization [8, 9] for hidden layer weight vectors. To examine its feasibility, we implement it in a large-scale DNN-HMM speech recognizer and conduct various evaluation experiments in a difficult task: TED Talks.

Quite recently, a gLasso-based idea similar to ours was reported [10]. However, we independently studied it and elaborated the characteristics of its scheme in large-scale experiments. In Section 4, we will discuss the relationship between our work and related works, including [10].



Fig. 1. Schematic explanation of outgoing/incoming weight vector grouping.

2. DEEP NEURAL NETWORK TRAINING ASSOCIATED WITH NODE SELECTION

2.1. Computational Procedures in DNNs

We consider a standard *L*-layer DNN, where the 0th layer is an input layer, the *L*-th layer is an output layer, and the 1st through (L-1)-th layers are hidden layers. In this network, the *l*-th layer consists of N_l nodes, each of which is associated with a bias coefficient. The nodes are also fully connected among the adjacent layers, and each connection is associated with a weight parameter.

When a vector pattern is given to the input layer, the network feeds the data from the input layer to the output layer in the following layer-by-layer manner:

$$\mathbf{a}^{l} = \mathbf{W}^{l} \mathbf{z}^{l-1} + \mathbf{b}^{l}, \text{ and } \mathbf{z}^{l} = \sigma(\mathbf{a}^{l}), \tag{1}$$

where \mathbf{W}^{l} is the weight matrix between the (l-1)-th and l-th layers, \mathbf{b}^{l} is the bias vector of the l-th layer, \mathbf{z}^{l-1} is the output vector from the (l-1)-th layer, \mathbf{a}^{l} is the activation vector in the l-th layer, and $\sigma()$ is an activation function. Here, the *i*-th row and *j*-th column element of \mathbf{W}^{l} are w_{ij}^{l} , which is the weight parameter between the *j*-th node of the (l-1)-th layer and the *i*-th node of the l-th layer; the vectors are detailed as follows: $\mathbf{a}^{l} = [a_{1}^{l} \cdots a_{i}^{l} \cdots a_{N_{l}}^{l}]$, $\mathbf{z}^{l-1} = [z_{1}^{l-1} \cdots z_{j}^{l-1} \cdots z_{N_{l-1}}^{l-1}]$ and $\mathbf{b}^{l} = [b_{1}^{l} \cdots b_{i}^{l} \cdots b_{N_{l}}^{l}]$, where a_{i}^{l}, z_{j}^{l-1} , and b_{i}^{l} are respectively the aggregated input to the *i*-th node of the (l-1)-th layer, and the bias of the *i*-th node of the *l*-th layer.

2.2. Node Selection without Activation Changes

Figure 1 illustrates an example situation of the node connections between the (l-1)-th layer (of the 2 nodes) and *l*-th layer (of the 2 nodes). As shown on its left and right sides, we can basically group the connections in two ways: 1) *outgoing* bundle or 2) *incoming* bundle. Based on these groupings, we further group their corresponding weights in outgoing weight vectors (\mathbf{w}_1^l and \mathbf{w}_2^l) or incoming weight vectors ($\tilde{\mathbf{w}}_1^l$ and $\tilde{\mathbf{w}}_2^l$). Here, the outgoing weight vectors are column

This work was supported in part by JSPS Grants-in-Aid for Scientific Research No. 26280063, MEXT-Supported Program "Driver-in-the-Loop," and Grant-in-Aid for JSPS Fellows.

vectors, which form the columns of weight matrix \mathbf{W}^{l} ; the incoming weight vectors form the rows of \mathbf{W}^{l} .

In the figure, for example, if the norm of \mathbf{w}_2^l is nearly zero and the norm for \mathbf{w}_1^l exceeds zero, \mathbf{a}^l is essentially based only on z_1^{l-1} , i.e., the outputs from the 1st (ν_1^{l-1}) node in the (l-1)-th layer. In this situation, even if we remove node ν_2^{l-1} , \mathbf{a}^l has no large changes and retains the DNN performances that were gained before node pruning. On the other hand, again for example, if the norm of $\mathbf{\tilde{w}}_2^l$ is nearly zero and the norm for $\mathbf{\tilde{w}}_1^l$ exceeds zero, $\mathbf{a}^{(l+1)}$ is essentially based on z_1^l and $b_2^l (\approx z_2^l)$. In this situation, even if we prune node ν_2^l , $\mathbf{a}^{(l+1)}$ has no large changes and the DNN performances accordingly remain almost the same. However, because the node is pruned based on its weight vector norm, its bias remains as a constant. Therefore, we have to remember here to shift the bias for the pruned node to the upper layer nodes.

As above, regardless of the weight vector grouping, if we can control the values of the weight vector norm in the DNN training, or more precisely, selectively minimize the norm values, we can also achieve a mechanism that automatically selects nodes without changes in the performances of the trained DNNs.

2.3. DNN Training with gLasso and L2 Regularizers

2.3.1. Definition of regularized loss

The node selection in 2.2 is naturally expected to be done so that only the salient nodes that produce useful outputs for increasing the DNN's performances are retained. If DNN training is done to concurrently minimize a loss function that reflects the DNN's classification errors and a function that reflects the norm values of the outgoing/incoming weight vectors, the unnecessary nodes, which produce useless outputs (for the DNN performances), are automatically disabled by their corresponding small norm values and only the salient nodes will be retained.

To implement such training, we adopt gLasso regularization for the outgoing or incoming weight vectors in standard DNN training using Cross Entropy (CE) loss. Focusing on the outgoing weight vector case, we formalize our proposed training scheme below. The resulting formalization is basically common to that in the incoming weight vector case. Note that the target layer for grouping is different between types of grouping.

Our training scheme starts by defining regularized loss L() as follows:

$$L(\mathbf{\Lambda}) = E(\mathbf{\Lambda}) + \alpha \hat{R}_{\text{group}}(\{\mathbf{W}^l\}_{l=2}^L) + \beta \hat{R}_{L2}(\mathbf{W}^1, \{\mathbf{b}^l\}_{l=1}^L), \quad (2)$$

where Λ is a set of trainable parameters of DNN, i.e., $\Lambda = \{\{\mathbf{W}^l, \mathbf{b}^l\}_{l=1}^L\}, E()$ is the CE loss, $\hat{R}_{group}()$ $\hat{R}_{L2}()$ are respectively the gLasso and L2 regularizers, and α and β are positive constants that control the effects of their corresponding regularizers. Here the gLasso regularizer is specified as

$$\hat{R}_{\text{group}}(\{\mathbf{W}^l\}_{l=2}^L) = \sum_{l=2}^L R_{\text{group}}(\mathbf{W}^l), \qquad (3)$$

where $R_{\text{group}}(\mathbf{W}^l) = \sum_{j=1}^{N_{l-1}} \|\mathbf{w}_j^l\|$, and the L2 regularizer is specified as

$$\hat{R}_{L2}(\mathbf{W}^{1}, \{\mathbf{b}^{l}\}_{l=1}^{L}) = R_{L2}(\mathbf{W}^{1}) + \sum_{l=1}^{L} R_{L2}(\mathbf{b}^{l}), \qquad (4)$$

where $R_{L2}(\mathbf{W}^l) = \frac{1}{2} \|\mathbf{W}^l\|^2$ and $R_{L2}(\mathbf{b}^l) = \frac{1}{2} \|\mathbf{b}^l\|^2$.

2.3.2. Effects of gLasso regularization

The status of Λ that corresponds to the minimum of $L(\Lambda)$ will provide high classification power with selected network nodes. To find that status, we minimize $L(\Lambda)$ using the standard gradient-based optimization procedure.

In the above gradient-based loss minimization for $L(\Lambda)$, the gradient terms of the regularizers play a key role in accelerating the minimization. In the L2 regularizer case, the gradient of $R_{L2}()$ is given as follows, for example, in terms of \mathbf{w}_{i}^{l} :

$$\frac{\partial R_{L2}(\mathbf{W}^l)}{\partial \mathbf{w}_j^l} = \mathbf{w}_j^l.$$
(5)

In contrast, the gradient of the gLasso regularizer becomes

$$\frac{\partial R_{\text{group}}(\mathbf{W}^l)}{\partial \mathbf{w}_j^l} = \frac{\mathbf{w}_j^l}{\|\mathbf{w}_j^l\|}.$$
(6)

Eq. (6) clearly indicates the effects of gLasso regularization; the ("outgoing" in this case) weight vector rapidly converges to zero when its norm is small. As a result, the node outputs, which are less useful for the minimization of $L(\Lambda)$, are disabled in an early stage of the loss minimization, and the retained node outputs will be dominantly used for the loss minimization.

2.4. Node Selection Procedure

If the weight vector norm becomes sufficiently small after the above loss minimization, we can simply prune its corresponding nodes using a threshold value (θ). The node selection procedure is summarized as follows: 1) calculate the gLasso₀ norm (column norm) or the gLasso₁ norm (row norm) for all of the hidden layer nodes, 2) prune the nodes whose norm values are less than θ , and 3) shift the bias for the pruned node to the upper layer nodes, especially in the gLasso₁ case. (See 3.2 for the definitions of gLasso₀ and gLasso₁.)

3. EXPERIMENTS

3.1. Settings

To evaluate our proposed DNN training and node selection method, we conducted speech recognition experiments using our DNN-HMM hybrid speech recognizer over a TED Talks corpus that consists of lecture speech data spoken by 793 speakers. For the evaluation experiments, we divided the data into the following three sets: training data by 774 speakers, validation data by 8 speakers, and testing data by 11 speakers. The total length of the training data was about 62 hours, the number of senone classes was 3,944, and the vocabulary size was 179, 604.

We represented the input speech as a series of 11 acoustic feature vectors, each of which was of 39 dimensions and consisted of MFCCs, log power, their first and second derivatives. The dimension of the concatenated input vector was 429.

In our DNN-HMM hybrid speech recognizer, the front-end DNN part was the 6-layer Multi-Layer Perceptron (L = 6 in 2.1) with a sigmoid activation function, and the post-end HMM part used a context-dependent acoustic model and a 3-gram language model. We initialized the DNN part by the RBM pre-training [11]. Independently of the DNN part, we developed a HMM part using Boosted MMI training [12].

The main target of our DNN training was to increase the senone classification accuracy (SCA), whose categories were indicated by the DNN outputs. However, we also evaluated the training results using the word error rate (WER) obtained by the post-end HMM part of the hybrid recognizer.

Regularizer	SCAB	WERB	SCAA	WERA
gLasso _O	46.2	18.7	46.2	18.8
gLasso _I	46.1	18.2	46.1	18.2
L2o	45.9	18.6	22.8	42.6
$L2_{I}$	45.9	18.6	8.2	100.0

Table 1. Achieved classification performances (%) in DNN₂₀₄₈

3.2. Evaluation procedures

We ran gradient-based training for minimizing $L(\Lambda)$ in (2) several times, along with changing such hyper-parameters as α and β in (2). Nevertheless, for simplicity, we restricted β to $\beta = 0.1\alpha$ or $\beta = 0$. In every training run, we repeated twenty epochs, in each of which the whole set of training samples was used only once. Among the different trained statuses of Λ , we selected the one that produced the highest SCA value over the validation data, evaluated the SCA values for the selected status of Λ over the testing data, and measured the WER values over the testing data using the selected DNN in the hybrid recognizer.

In addition to experiments with our proposed training associated with the gLasso regularizer, we also tested (for comparison purposes) the training using the L2 regularizer for all of the weight matrixes and all of the bias vectors. Then, the loss in (2) was replaced in this additional case by

$$L(\mathbf{\Lambda}) = E(\mathbf{\Lambda}) + \beta \sum_{l=1}^{L} \left\{ R_{L2}(\mathbf{W}^{l}) + R_{L2}(\mathbf{b}^{l}) \right\}.$$
 (7)

To elaborate the effects of our proposed training method, we also tested two different DNN sizes: one with 2048 nodes for every hidden layer (DNN₂₀₄₈) and another with 4096 nodes for every hidden layer (DNN₄₀₉₆). For both DNN₂₀₄₈ and DNN₄₀₉₆, we evaluated the following four types of training: 1) with the gLasso regularizer for the outgoing hidden layer weight vectors (gLasso₀), 2) with the gLasso regularizer for the incoming hidden layer weight vectors (gLasso_I), $\overline{3}$) with the L2 regularizer for the outgoing hidden layer weight vectors (L2₀), and 4) with the L2 regularizer for the incoming hidden layer weight vectors (L2_I). Particularly in the L2₀ and L2_I cases, we conducted common L2-regularized training (which corresponds to the training using (7)) but selected different nodes, based on the differences in grouping. Moreover, as stated in 2.3, we used in all four cases the L2 regularizer for all of the bias vectors; We also used the L2 regularizer for the input layer weights in the $gLasso_O$ case, and for the output layer weights in the $gLasso_I$ case.

For convenience, we collectively refer to both the $gLasso_0$ and $gLasso_1$ training cases as gLasso and L2 for both the $L2_0$ and $L2_1$ cases.

3.3. Results

3.3.1. Fundamental performances

In Table 1, we summarize the SCA and WER values that we achieved for DNN₂₀₄₈ in the four training cases: gLasso₀, gLasso₁, L2₀, and L2₁. SCA_B and WER_B represent the SCA and WER values (%), both obtained before the node selection; the SCA and WER values were obtained after the node selection for SCA_A and WER_A.

First, we focus on the results in the SCA_B and WER_B columns. Regardless of the weight vector grouping type and the regularizer type, DNN training achieved such accurate classification performances as an SCA of about 46% and a WER less than 19%. They also show that the gLasso procedures achieved the performance that were competitive with the L2 procedures, which are widely used in the neural network training.



Fig. 2. Histogram of outgoing weight vector norm values in $gLasso_0$ or $L2_0$ cases. Training was done for DNN₂₀₄₈. Top: $gLasso_0$, Bottom: $L2_0$.

Next, we observe the characteristics of the weight vector norm values that were obtained by the training. Figure 2 shows a histogram of the outgoing weight vector norm values in the gLasso₀ or L2₀ cases for DNN₂₀₄₈. The horizontal axis indicates the weight vector norm value, and the vertical axis indicates the frequency of its corresponding norm value. From this figure, we obtained the following findings: 1) The weight vector norm values by the gLasso₀ cases were clearly separated into two groups, suggesting the ease with which the threshold (θ) can be set for the node selection between them; 2) The weight vector norm values reduced by the gLassoo procedure were rather small (at most slightly larger than 10^{-4}). This distinct reduction proves that the gLassoo process successfully reduced the norm values of some selected weight vectors. 3) The weight vector norm values by the L2₀ case were concentrated in the region between 10^0 and 10^1 . This implies that the L2 procedure failed to reduce the weight vector norm values.

Due to space limitations, we omit the introduction of the results in the $gLasso_I$ and $L2_I$ cases. However, we obtained norm value distribution that closely resembled that in Figure 2.

We return to Table 1 and compare the results in the SCA_B and WER_B columns with those in the SCA_A and WER_A columns.

We executed node selection by setting θ to 10^{-2} , which is clearly far from the norm values in either of the two separate regions generated by the gLasso training. Using this threshold, we pruned 3161 nodes from the five hidden layers in the gLasso₀ case (the number of retained hidden layer nodes = 7079) and 3368 nodes in the gLasso₁ (number of retained hidden layer nodes = 6872) case. The performances gained after the node selection (in the SCA_A and WER_A columns) are basically the same with those gained before the node selection (in the SCA_B and WER_B columns). The comparison here clearly proves that, regardless of the type of weight vector grouping, gLasso training sorted out the weight vectors by selectively reducing the norm values for some weight vectors and allowed node selection simply based on thresholding.

As Figure 2 shows, pruning a certain amount of nodes is not easy from L2-based DNNs by setting a threshold. For comparison, we removed (in ascending order) the small-norm nodes from the DNN produced by the common L2-regularized training in the L_{20} case until the number of removed nodes became identical to the gLasso₀ case: 3161; similarly, in the L2₁ case, we removed 3368 nodes from the L2-based DNN, which was the same with the network used in the L2₀ case. The scores in the SCA_A and WER_A columns of Table 1 were obtained using these size-reduced L2-based DNNs. They are much worse than those gained before the node selection and clearly show that both the L2₀ and L2₁ cases failed to sort out the weight vectors (as a result, their corresponding nodes).

To further analyze the effects of regularizer selection, we observed the changes in the SCA scores along with gradually removing hidden layer nodes, by 100 nodes, in increasing order of their norm values. Figure 3 shows the SCA scores, each of which is a function of the number of removed nodes. In the figure, the horizontal axis



Fig. 3. Effects of regularizer selection in node selection in DNN₂₀₄₈.



Fig. 4. Number of disposal nodes for every hidden layer as a function of training epoch in gLasso₀-based DNN₂₀₄₈.

indicates the number of removed nodes; the SCA scores for the vertical axis. If the nodes do not exist in the small norm value region, the SCA scores gradually decrease as nodes are removed. On the other hand, if the nodes exist in the small norm value region, the SCA values remain high and later decrease. The curves in the figure clearly prove that the L2 procedure corresponds to the former situation and the gLasso procedure achieves the latter desirable situation.

3.3.2. Properties of gLasso procedure

When introducing gLasso's training scheme in 2.3, we stated that the disposable nodes are disabled in some early training stage. To investigate this point, we observed the number of disposable hidden layer nodes for every hidden layer as a function of the training epoch. Here we define a disposable node as one whose norm value is smaller than 10^{-2} . Figure 4 illustrates the observed numbers for all five hidden layers (1st through 5th) of the gLasso₀-based DNN₂₀₄₈. We found the following: 1) Node selection was stably realized with the training progress. 2) It occurred in such comparatively early epochs as the 1st through 5th epochs. 3) It dominantly occurred in such higher hidden layers as the 3rd and 4th.

The results in Figure 4 suggest that the gLasso procedure shows a trend that mainly reduces the weight vector norm values in the higher hidden layers. To scrutinize this trend, in both the DNN_{2048} and DNN_{4096} cases, we pruned the disposable nodes using the 10^{-2} threshold and observed the number of selected (retained) nodes for every hidden layer. Figure 5 illustrates those numbers. Regardless of the weight vector grouping (outgoing vs. incoming) or the DNN size (DNN_{2048} vs. DNN_{4096}), we found that the gLasso procedure suppressed the weight vector norm values, specifically in the 3rd and 4th hidden layers, and achieved a size-reduced and reshaped network structure that was shared in all of the tested cases of gLasso₀, gLasso₁, DNN_{2048} , and DNN_{4096} .

4. RELATED WORK

An orthodox approach for the size reduction of a large matrix is to use Singular Value Decomposition. A large weight matrix in DNN is decomposed into small matrixes and re-expressed (in the sense of low-rank approximation) by replacing the small eigenvalues of the decomposed diagonal matrix by zero [4, 5]. This approach allows



Fig. 5. Number of finally retained hidden layer nodes in gLasso₀ and gLasso₁ cases.

parameter reduction without serious degradation of the original discriminative power of the network. However, it does not reduce the number of hidden layer nodes, and therefore it is probably insufficient to achieve a desirable amount of network parameter reduction, especially in the case of using the network with the huge number of hidden layer nodes.

For directly pruning the hidden layer nodes, the L1 norm of the weight vectors was used as an *important* function. It resembles the weight-grouping-based norm that we use in our proposed training procedure [7]. However, because such mechanism of reducing the norms in the training procedure as ours was not implemented, DNN restructuring had to rely on additional retraining results.

To find a desirable DNN structure, the Genetic Algorithm (GA) has been applied [13]. It has high potential for generating a target structure, but it often needs an enormous amount of computational resources.

Compared with the above related studies, our proposed training that embeds the gLasso regularizer is characterized by two main advantages: 1) Even if no additional training is conducted, it can easily produce node-pruned DNNs without performance degradation, and 2) its computation time is almost the same as conventional L2 regularization.

Motivated by the gLasso procedure defined in linear regression [8], approaches similar to ours were proposed [10, 14]. In [10], the gLasso regularizer was defined for the outgoing vector grouping of DNN and its effect was evaluated in various small- and middle-sized tasks that classify fixed-dimensional patterns. In [14], the gLasso regularizer was defined for the incoming vector grouping¹ of neural network-based language models and its effect was evaluated for shallow networks. Compared to these work, we studied in this paper both the incoming and outgoing vector groupings and also elaborated our proposed scheme in a large-scale speech recognition environment using larger network structures.

5. CONCLUSION

We investigated the feasibility of the gLasso procedure, i.e., DNN training associated with the gLasso regularizer, for hidden layer weight vectors in a large-scale speech recognition task with the DNN-HMM hybrid speech recognizer in a TED Talks task. We elaborated the nature of gLasso training by applying it to both outgoing and incoming weight vector groupings and compared the gLasso and L2 procedures, where the L2 regularizer was used for all of the weight matrixes and all of the bias vectors. From the experiment results, we clearly demonstrated that our gLasso procedure automatically (without additional training) disabled less useful hidden layer nodes without degradation in DNN classification performance and successfully produced a small, necessary, and sufficient DNN structure.

¹The usage of bias term is different from our formalization.

6. REFERENCES

- [1] G. Hinton, L. Deng, D. Yu, G. Dahl, A. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. Sainath, and B. Kingsbury, "Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups," Signal Processing Magazine, IEEE, vol. 29, no. 6, pp. 8297, 2012.
- [2] T. Ochiai, S. Matsuda, X. Lu, C. Hori, and S. Katagiri, "Speaker adaptive training using deep neural networks," in Proc. ICASSP, 2014, pp. 6349-6353.
- [3] T. Ochiai, S. Matsuda, H. Watanabe, X. Lu, C. Hori, and S. Katagiri, "Speaker adaptive training for deep neural networks embedding linear transformation networks," in Proc. ICASSP, 2015, pp. 4605-4609.
- [4] J. Xue, J. Li, and Y. Gong, "Restructuring of deep neural network acoustic models with singular value decomposition," in Proc. INTERSPPECH, 2013, pp. 2365-2369.
- [5] J. Xue, J. Li, D. Yu, M. Seltzer, and Y. Gong, "Singular value decomposition based low-footprint speaker adaptation and personalization for deep neural network," in Proc. ICASSP, 2014, pp. 6359-6363.
- [6] T. Ochiai, S. Matsuda, H. Watanabe, X. Lu, H. Kawai, and S. Katagiri, "Bottleneck linear transformation network adaptation for speaker adaptive training-based hybrid DNN-HMM speech recognizer," in Proc. ICASSP, 2016, pp. 5015-5019.
- [7] T. He, Y. Fan, Y. Qian, T. Tan, and Kai Yu, "Reshaping deep neural network for fast decoding by node-pruning," in Proc. ICASSP, 2014, pp. 245-249.
- [8] M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," Journal of the Royal Statistical Society: Series B (Statistical Methodology), vol. 68, no. 1, pp. 49-67, 2006.
- [9] N. Simon, J. Friedman, T. Hastie, and R. Tibshirani, "A sparsegroup lasso" Journal of Computational and Graphical Statistics, vol. 22, no. 2, pp. 231-245, 2013.
- [10] S. Scardapane, Danilo Comminiello, Amir Hussain, and Aurelio Uncini, "Group Sparse Regularization for Deep Neural Networks," arXiv:1607.00485v1 (2016).
- [11] G. E. Hinton, S. Osindero, and Y. Teh, "A fast learning algorithm for deep belief nets," Neural Computation, vol. 18, pp. 1527-1554, 2006.
- [12] D. Povey, D. Kanevsky, B. Kingsbury, B. Ramabhadran, G. Saon, and K. Visweswariah, "Boosted MMI for model and feature space discriminative training," in Proc. ICASSP, 2008, pp. 4057-4060.
- [13] T. Shinozaki and S. Watanabe, "Structure discovery of deep neural network based on evolutionary algorithms," in Proc. ICASSP, 2015, pp. 4979-4983.
- [14] K. Murray and D. Chiang, "Auto-sizing neural networks: With applications to n-gram language models," in Proc. EMNLP, 2015, pp. 908916.