

# SelectScale: Mining More Patterns from Images via Selective and Soft Dropout

Zhengsu Chen<sup>1</sup>, Jianwei Niu<sup>1,2,3</sup>, Xuefeng Liu<sup>1\*</sup> and Shaojie Tang<sup>4</sup>

<sup>1</sup>Beihang University

<sup>2</sup>Hangzhou Innovation Institute of Beihang University

<sup>3</sup>Zhengzhou University

<sup>4</sup>University of Texas at Dallas

{danczs, niujianwei, liu\_xuefeng}@buaa.edu.com, tangshaojie@gmail.com

## Abstract

Convolutional neural networks (CNNs) have achieved remarkable success in image recognition. Although the internal patterns of the input images are effectively learned by the CNNs, these patterns only constitute a small proportion of useful patterns contained in the input images. This can be attributed to the fact that the CNNs will stop learning if the learned patterns are enough to make a correct classification. Network regularization methods like dropout and SpatialDropout can ease this problem. During training, they randomly drop the features. These dropout methods, in essence, change the patterns learned by the networks, and in turn, forces the networks to learn other patterns to make the correct classification. However, the above methods have an important drawback. Randomly dropping features is generally inefficient and can introduce unnecessary noise. To tackle this problem, we propose SelectScale. Instead of randomly dropping units, SelectScale selects the important features in networks and adjusts them during training. Using SelectScale, we improve the performance of CNNs on CIFAR and ImageNet.

## 1 Introduction

Recently, deep neural networks have enabled breakthroughs on computer vision and natural language processing tasks [Girshick *et al.*, 2014; Long *et al.*, 2015; Graves *et al.*, 2013]. With specially designed architecture, convolutional neural networks (CNNs) have been widely used on image classification [Szegedy *et al.*, 2015; He *et al.*, 2016a], object detection [Ren *et al.*, 2015] and semantic segmentation [Long *et al.*, 2015]. Since AlexNet [Krizhevsky *et al.*, 2012] significantly improved the classification performance on ImageNet, many new deep convolutional neural network architectures (VGG [Simonyan and Zisserman, 2014], Inception [Szegedy *et al.*, 2015; 2016], ResNet [He *et al.*, 2016a], DenseNet [Huang *et al.*, 2016a]) are proposed. CNNs extract the features of the input data automatically and the classifier makes decisions according to the features.

\*Contact Author

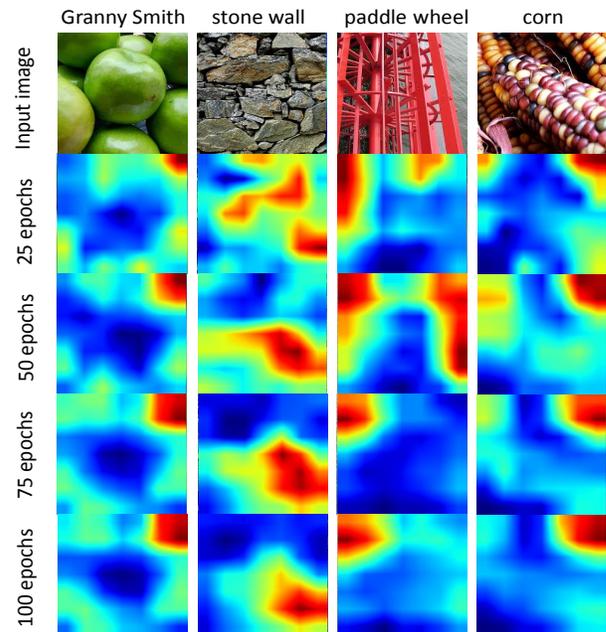


Figure 1: The class activation maps of training examples for different training epochs. During training, the networks may only utilize a small subset of all useful patterns in images. Even if they grasp rich patterns in the beginning, they can lose some of them during training and overfit to the remaining patterns.

However, CNNs usually can not utilize all of the useful patterns in the input data for image classification. If they have learned enough patterns for classification, they will stop learning from the samples. Figure 1 shows the class activation maps (CAM) [Zhou *et al.*, 2016] of training images for ResNet-50 on ImageNet. As can be seen, during training the network only learns a subset of patterns from these samples. However, there still exist many patterns yet to be observed by the CNNs. These patterns, if added into the CNNs, will increase the robustness and accuracy of the networks.

We observed that some regularization techniques can ease the problem above. The most popular one is dropout [Srivastava *et al.*, 2014]. Dropout randomly suppresses the outputs of neural units in networks. For a well-learned sample,

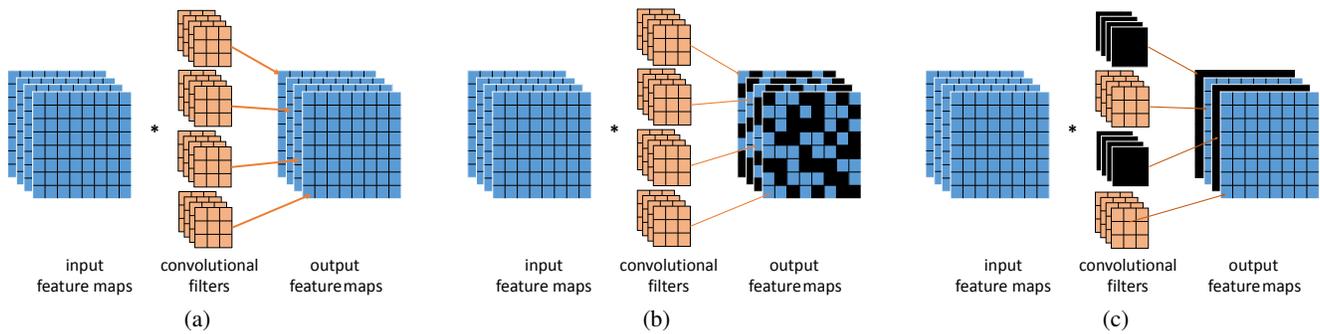


Figure 2: (a) A standard  $3 \times 3$  convolutional layer. (b) Standard dropout for outputs. The dropped elements are shown in black. The retaining rate is 0.5. (c) SpatialDropout for outputs. The dropped feature maps are shown in black. It is equal to dropping the corresponding filters. The retaining rate is 0.5.

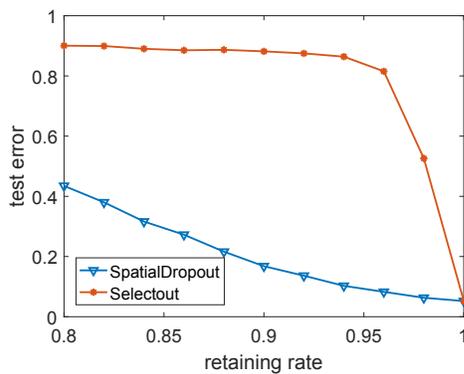


Figure 3: The test error for WRN-20-4 with different data drop methods on CIFAR-10 for inference.

dropout will change its features. If the learned patterns are destroyed by dropout, the networks need to learn new patterns to recognize the input image. Thus a network with dropout generally outperforms the original networks.

However, dropout suppresses the output randomly. It may not drop the data as we want. First, it may drop the features of unlearned patterns. Under this circumstance, it has no contribution to learning new patterns but introduces unnecessary noise to the networks. Second, dropout may drop too many important features. The remaining information may be not enough for recognition, which may mislead the networks and make the networks unstable.

To address this problem, we propose Selectout. Selectout suppresses the neural units selectively instead of dropping them randomly. During training, we design a fast and parameter-free method to select the features that are important for recognition. Selectout changes these important features during training, therefore the networks need to learn new useful patterns to classify it. For Selectout, the dropped features are under control. First, it only drops the features that are important for the networks, which avoids introducing unnecessary noise. Second, Selectout can determine how many

important features to drop. Thus, the features will not be seriously destroyed by the drop method. Selectout is a more effective data drop strategy.

Since SpatialDropout [Tompson *et al.*, 2015] is more effective than dropout in dropping the features, we construct Selectout with SpatialDropout. SpatialDropout directly drops feature maps instead of dropping neural units. The difference between standard dropout and SpatialDropout is shown in Figure 2(b) and 2(c). Based on SpatialDropout, Selectout needs to select the important feature maps in a convolutional layer.

Selectout is a very strong regularization method because it drops the important features in the networks. To roughly show the regularization capacity, we train a network without drop method and apply data drop methods to it during test on CIFAR-10. As shown in Figure 3, Selectout increases the test error more rapidly than SpatialDropout as the retaining rate decreases. The feature maps selected by Selectout are much more important than the feature maps that are randomly selected by SpatialDropout. Only dropping 6% of the feature maps by Selectout almost disables the network. In summary, our selection approach for important feature maps is very effective and Selectout is a very strong regularization method.

The regularization of Selectout based on SpatialDropout can be too strong for the networks. The networks are usually very sensitive to the data retaining rate and hard to be tuned. To overcome this problem, we propose SelectScale. SelectScale is based on SpatialScale. Unlike SpatialDropout, SpatialScale scales the feature maps with random weights rather than directly sets some of them to zeros. SpatialScale is similar to SpatialDropout but is milder. It is more stable and easier to be tuned than SpatialDropout. SelectScale is the selecting drop method based on SpatialScale. It only scales the selected feature maps. SelectScale is the main method used in this paper.

SelectScale based on SpatialScale is tested on ResNets, WRNs, ResNeXt and non-residual networks in this paper. It is observed that SelectScale can consistently improve the networks and outperform standard dropout, SpatialDropout, and SpatialScale.

To summarize, our main contributions are as follows:

- We propose a novel data drop method that drops the data selectively rather than randomly, which encourages the networks to learn more patterns.
- We present a fast and parameter-free method to rank the feature maps during data drop and propose a mild way of dropping data.
- We conduct detailed experiments to qualitatively and quantitatively demonstrate that our data drop method can significantly improve the network performance.

## 2 Methods

In this section, we first review data drop methods and construct Selectout based on SpatialDropout (section 2.1). Since Selectout is too strong for many networks, we introduce our mild data drop method: SelectScale (section 2.2). Finally, we detail the function of selecting feature maps when dropping data (section 2.3).

### 2.1 Selectout Based on SpatialDropout

Given an input image,  $\mathbf{x}$  is the input of a convolutional layer for a CNN.  $c$  is the filter number of this layer.  $Y$  are the output feature maps and  $W = \{\mathbf{w}_i | 1 \leq i \leq c\}$  are the weights of the filters in this layer. We have:

$$\begin{aligned} \mathbf{y}_i &= f(\mathbf{w}_i * \mathbf{x} + b_i) \\ Y &= \{\mathbf{y}_i | 1 \leq i \leq c\} \end{aligned} \quad (1)$$

where  $i$  indexes the filters and the output feature maps in the layer.  $f$  is the activation function.  $b_i$  is the bias.

Using standard dropout, the outputs will be masked by a Bernoulli distribution:

$$\begin{aligned} \mathbf{y}_i^{dropout} &= \mathbf{r}_i \odot \mathbf{y}_i / p \\ \mathbf{r}_i &\sim \text{Bernoulli}(p) \\ p &\in [0, 1] \\ Y_{dropout} &= \{\mathbf{y}_i^{dropout} | 1 \leq i \leq c\} \end{aligned} \quad (2)$$

where  $p$  is the data retaining rate. The shape of  $\mathbf{r}_i$  is the same as  $\mathbf{y}_i$ . To compensate the change of data distribution by dropout, the outputs are scaled by  $1/p$  during training [Abadi *et al.*, 2016]. An alternative method is scaling the outputs by  $p$  during testing instead [Srivastava *et al.*, 2014].

Using SpatialDropout, the outputs:

$$\begin{aligned} \mathbf{y}_i^{SpatialDropout} &= r_i \mathbf{y}_i / p \\ r_i &\sim \text{Bernoulli}(p) \\ p &\in [0, 1] \\ Y_{SpatialDropout} &= \{\mathbf{y}_i^{SpatialDropout} | 1 \leq i \leq c\} \end{aligned} \quad (3)$$

where  $r_i$  is the weight to determinate whether to drop the feature map  $\mathbf{y}_i$ . Like standard dropout, data after SpatialDropout is also scaled by  $1/p$  during training and data drop will be removed during testing.

SpatialDropout regards the convolution filter rather than the neural unit as the basic data drop unit. In convolutional neural networks, the units in the same feature map share the same weights. Thus SpatialDropout directly drops feature maps to regularize the networks as shown in Figure 2(c). This

is the main difference between standard dropout and SpatialDropout.

Based on SpatialDropout, Selectout is:

$$\begin{aligned} \mathbf{y}_i^{Selectout} &= r_i \mathbf{y}_i / p \\ \mathbf{r} &= S(Y, t, p) \\ r_i &\in \{0, 1\} \\ Y_{Selectout} &= \{\mathbf{y}_i^{Selectout} | 1 \leq i \leq c\} \end{aligned} \quad (4)$$

where  $S$  is the selecting function, which determinates how to drop the feature maps.  $S$  first generates a score for each feature map and ranks the feature maps with the scores. Then  $S$  selects some of the feature maps according to  $t$  and  $p$ .  $t$  is the top rate and  $p$  is the data retaining rate. For example, using  $t = 0.2$  and  $p = 0.9$  indicates that 10% of the feature maps are dropped and these dropped feature maps are randomly selected from the top 20% feature maps.  $t$  and  $p$  are the hyperparameters that can be tuned.  $S$  is detailed in section 2.3.

### 2.2 SelectScale Based on SpatialScale

SpatialDropout directly suppresses the outputs of filters. Since SpatialDropout with Selectout will supply stronger regularization, a very small drop rate (e.g. 0.02) can be enough for some networks. However, sometimes the feature maps in convolutional layers can be too few to reveal the difference between very small drop rates. The regularization supplied by Selectout may be unstable and the hyperparameters can be hard to be tuned. We, therefore, propose SpatialScale, which randomly scales the feature maps. Using SpatialScale:

$$\begin{aligned} \mathbf{y}_i^{SpatialScale} &= r_i \mathbf{y}_i \\ r_i &\sim \text{Uniform}(1 - q, 1 + q) \\ q &\in [0, 1] \\ Y_{SpatialScale} &= \{\mathbf{y}_i^{SpatialScale} | 1 \leq i \leq c\} \end{aligned} \quad (5)$$

The outputs do not need to be scaled, because the mean value of  $r_i$  is 1. During testing, SpatialScale is removed.

Both SpatialDropout and SpatialScale focus on changing the outputs to regularize the convolutional layers. SpatialDropout directly suppresses the outputs. SpatialScale randomly scales the outputs instead. SpatialScale is milder and more stable than SpatialDropout. It changes the output data modestly.

Based on SpatialScale, SelectScale is:

$$\begin{aligned} \mathbf{y}_i^{SelectScale} &= r_i \mathbf{y}_i \\ r_i &= (1 - v_i)(u_i - 1) + 1 \\ u_i &\sim \text{Uniform}(1 - q, 1 + q) \\ q &\in [0, 1] \\ \mathbf{v} &= S(Y, t, 1 - t) \\ v_i &\in \{0, 1\} \\ Y_{SelectScale} &= \{\mathbf{y}_i^{SelectScale} | 1 \leq i \leq c\} \end{aligned} \quad (6)$$

SelectScale only scales the feature maps that are selected by function  $S$ . The selecting function is detailed in section 2.3.

### 2.3 Selecting Function $S$

The selecting function takes the feature maps ( $Y$ ), the top rate ( $t$ ) and the data retaining rate ( $p$ ) as the inputs. In function  $S$ , we first calculate the absolute values of the feature maps. For the feature maps after ReLU, this step can be removed. Then, the maximum value in each feature map is calculated and these values are regarded as the score of corresponding feature maps. Next, the feature maps are ranked by the scores and the feature maps with top  $t$  scores are selected as candidates. For Selectout, the retaining feature maps are randomly selected from the candidates according to the data retaining rate  $p$ . For SelectScale, all of the candidates will be selected ( $p = 1 - t$ ).

In selecting function  $S$ , the maximum value in each feature map is used as the ranking score. Great activation values usually mean that useful patterns are detected. The results in Figure 3 prove that the feature maps with greater maximum values do play more important roles in neural networks.

Many other methods for obtaining the scores are tested. For example, we apply an average pooling layer before calculating the maximum value. Under these circumstances, the feature maps with greater average local activation values have higher scores. We test the average pooling layer with different pooling windows and strides. Global average pooling is also tested, which utilizes the average activation values of the feature maps as the scores. However, the performance gets worse when enlarging the pooling window and directly adopting the maximum values outperforms these methods.

## 3 Experiments

### 3.1 Datasets and Settings

CIFAR-10 and CIFAR-100 are colored  $32 \times 32$  natural image datasets that consist of 50,000 images in the training set and 10,000 images in the testing set. Following the common practice [He *et al.*, 2016a; Zagoruyko and Komodakis, 2016], these tiny images are padded with four pixels on each side and are randomly cropped with a  $32 \times 32$  window afterwards. If not specified, the networks for CIFAR are trained for 200 epochs with 0.1 initial learning rate. The learning rate drops by 0.2 at 60, 120 and 160 epochs following [Zagoruyko and Komodakis, 2016]. The batch size is 128 and the weight decay is 0.0005. The data drop method is applied to every convolutional layer for ResNets and WRNs. For ResNeXt, data drop is used before and after the  $3 \times 3$  group convolutions.

ILSVRC2012 is a subset of ImageNet [Russakovsky *et al.*, 2015] database. This dataset contains 1.3M images in the training set and 50K images in the validation set. The top-1 and top-5 error rates on the validation set are reported in this paper. On ImageNet, the standard color augmentation in [Russakovsky *et al.*, 2015] is utilized. Following the common practice [Szegedy *et al.*, 2016; He *et al.*, 2016b], horizontal flip, scale, and aspect ratio augmentation [Szegedy *et al.*, 2016] is utilized in the experiments. The batch size is 256 and the weight decay is 0.0001. The model is trained for 105 epochs with a 0.1 initial learning rate. The learning rate drops by 0.1 at 30, 60, 90 and 100 epochs. Data drop regularization is adopted for every convolutional layers in bottleneck resid-

ual blocks. Like [Szegedy *et al.*, 2016], the data drop methods are only utilized for the last two groups in ResNet-50.

We use the template in [He *et al.*, 2016a] to construct the ResNets on CIFAR. WRNs are constructed by widening the ResNets in this paper. WRN-20-4 represents the network with four times more filters in each layer than ResNet-20.

### 3.2 Different Top Rates and Retaining Rates for SelectScale

To reveal the effects of the two hyperparameters (top rate and retaining rate), SelectScale is tested with different top rates ( $t$ ) and retaining rates ( $q$ ) in this section. The experiments are conducted on CIFAR-100 with WRN-20-8. Fig. 4(a) shows results of different top rates with a fixed retaining rate ( $q = 0$ ). Fig. 4(a) presents the results of adopting different retaining rates with a fixed top rate ( $t = 0.2$ ). The results of the network without data drop and the best results of adopting other drop methods are also shown in the figures. As can be seen, SelectScale outperforms the best results of other methods in a wide range of top rates and retaining rates.

In Fig. 4(a), the network performs best when utilizing a quite small top rate (0.2). It indicates that regularizing the important feature maps is more effective than regularizing all the feature maps. Changing the important feature maps can encourage the network to learn more patterns, while changing the unimportant feature maps may introduce unnecessary noise during training.

### 3.3 Results on ResNets and WRNs

In this section, SelectScale is tested with ResNets and WRNs on CIFAR-10 and CIFAR-100. We apply the drop methods to every convolutional layer. The networks are tested using different retaining rates and the best results are shown in Table ?? and Table ??.

As shown in Table ??, SpatialDropout and SpatialScale consistently outperform standard dropout for all networks. When only regularizing the important feature maps by SelectScale, the regularization is more effective and the networks achieve better results. WRN-20-4 with SelectScale achieves a comparable result with WRN-20-8 using SpatialDropout. Note that WRN-20-8 employs about four times more parameters than WRN-20-4. When SelectScale is used for WRN-20-8, its performance is further improved.

On CIFAR-100, the improvement of SelectScale is more remarkable. As shown in Table ??, standard drop can only provide modest improvement. SpatialDropout and SpatialScale perform much better than standard dropout. SelectScale further reduces the test error by about 1% for wide residual networks. Note that the top rates for ResNet-110, WRN-20-4 and WRN-20-8 are 0.1, 0.2 and 0.1. In other words, SelectScale regularizes much fewer feature maps than SpatialScale and achieves better results.

### 3.4 SelectScale with Other Methods

In this section, more experiments are conducted to demonstrate that SelectScale does not conflict with new residual network architecture and other training methods. The methods utilized in the experiments are followings:

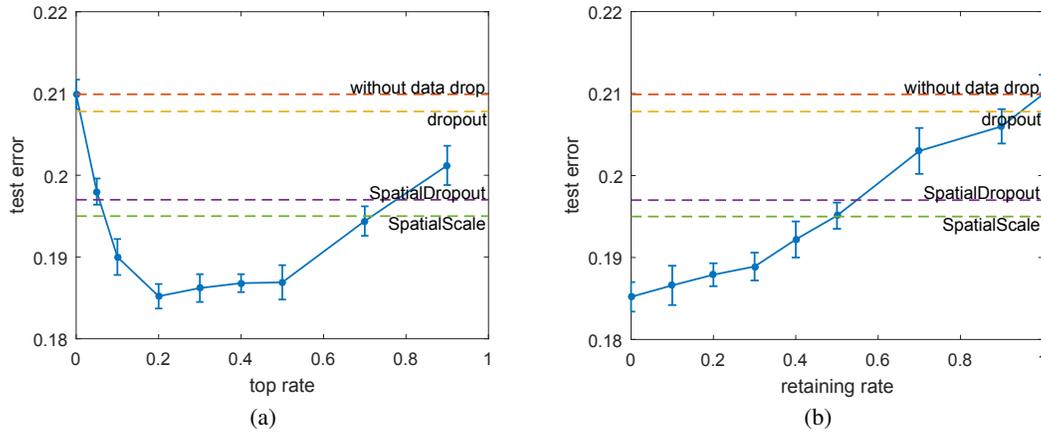


Figure 4: The test errors of different top rates and data retaining rates for WRN-20-8 on CIFAR-100. The best results of other methods are shown. (a) The test errors of the network with different top rates and a fixed retaining rate ( $q = 0$ ). (b) The test errors of the network with different retaining rates and a fixed top rate ( $t = 0.2$ ). Note that  $retainingrate = 0$  in SelectScale indicates that the top-ranked features are scaled by factors sampled from  $Uniform(0, 2)$ .

| Network    | without data drop(%) | standard dropout(%) | SpatialDropout(%) | SpatialScale(%) | SelectScale(%)  |
|------------|----------------------|---------------------|-------------------|-----------------|-----------------|
| ResNet-110 | $5.88 \pm 0.13$      | $5.78 \pm 0.11$     | $5.42 \pm 0.15$   | $5.40 \pm 0.06$ | $5.22 \pm 0.08$ |
| WRN-20-4   | $5.04 \pm 0.07$      | $4.84 \pm 0.06$     | $4.55 \pm 0.07$   | $4.64 \pm 0.12$ | $4.24 \pm 0.05$ |
| WRN-20-8   | $4.93 \pm 0.08$      | $4.37 \pm 0.07$     | $4.21 \pm 0.11$   | $4.14 \pm 0.05$ | $3.89 \pm 0.10$ |

Table 1: Test error of different CNNs on CIFAR-10. The data retaining rate ( $p$ ) of standard dropout and SpatialDropout is 0.9 for all tested networks. For SpatialScale, the retaining rate ( $q$ ) is 0.4. The top rates ( $t$ ) and data retaining rates ( $q$ ) of these three networks are 0.4, 0.2, 0.2 and 0.4, 0, 0.1. The results are the average of 5 runs.

- New Residual Architecture.** ResNeXt is a new residual network architecture that using group convolutions for bottleneck residual blocks [Xie *et al.*, 2016]. It introduces cardinality to residual networks and achieves better results than ResNets upon many visual recognition tasks.
- Cosine Annealing for Learning Rate.** Cosine annealing for learning rate is introduced by [Loshchilov and Hutter, 2016]. Instead of dropping the learning rate with a factor after each training stage, they drop the learning rate more smoothly by cosine annealing.
- Curriculum Dropout.** Curriculum dropout is introduced by [Morerio *et al.*, 2017]. In their opinion, the co-adaptations will not occur at the beginning of training, because the networks are initialized randomly. Thus they linearly decrease the data retaining rate during training.
- More Training Epochs.** Training the network for more epochs usually will improve the network performance on CIFAR-10 [Loshchilov and Hutter, 2016]. The training epochs are increased from 200 to 600 in our experiments.

The results on CIFAR-10 is shown in Table ?? . As can be

seen, because of introducing too much unnecessary noise, the test error is increased by standard dropout. SpatialDropout can improve the networks but SelectScale achieves the best result.

### 3.5 Results on ImageNet

ImageNet is a very large and complicated dataset and is very different from CIFAR. Many regularization methods that are useful in CIFAR can not improve the network performance on ImageNet [Huang *et al.*, 2016b; DeVries and Taylor, 2017]. In this section, we test our method on ImageNet. The tested network is ResNet-50. Drop methods are applied to the third and fourth groups after every BN [Ioffe and Szegedy, 2015] layer and the drop rate in the third group is decreased by a factor following [Ghiasi *et al.*, 2018]. The factor is 0.5 in our experiments. The data retaining rate for dropout and SpatialDropout is 0.9. The top rate for SelectScale is 0.1 and the data retaining rate is 0.2. Curriculum drop is utilized.

As shown in Table ??, standard dropout increases the test error. SpatialDropout almost makes no difference on the network performance. Both these two methods introduce too much noise to the network. SelectScale is more effective in adjusting the learned patterns and decreases the top-1 and top-5 error by 0.5% and 0.3% respectively.

| Network    | without data drop(%) | standard dropout(%) | SpatialDropout(%) | SpatialScale(%) | SelectScale(%) |
|------------|----------------------|---------------------|-------------------|-----------------|----------------|
| ResNet-110 | 26.16 ± 0.14         | 25.85 ± 0.24        | 24.12 ± 0.29      | 24.12 ± 0.19    | 23.48 ± 0.22   |
| WRN-20-4   | 22.90 ± 0.21         | 22.53 ± 0.23        | 21.24 ± 0.17      | 21.40 ± 0.08    | 20.20 ± 0.18   |
| WRN-20-8   | 20.99 ± 0.18         | 20.78 ± 0.19        | 19.70 ± 0.27      | 19.50 ± 0.23    | 18.52 ± 0.22   |

Table 2: Test error of different CNNs on CIFAR-100. The data retaining rate ( $p$ ) of standard dropout and SpatialDropout is 0.9 for all tested networks. For SpatialScale, the retaining rate ( $q$ ) is 0.4 except for WRN-20-8 (0.2). The top rates ( $t$ ) and data retaining rates ( $q$ ) of these three networks are 0.1, 0.1, 0.2 and 0.1, 0.1, 0.1. The results are the average of 5 runs.

| 200 ep. | 600 ep. | CA | CD | no drop | D | SD | SS | error (%) |
|---------|---------|----|----|---------|---|----|----|-----------|
| ✓       |         |    |    | ✓       |   |    |    | 4.19      |
| ✓       |         | ✓  |    | ✓       |   |    |    | 4.07      |
|         | ✓       | ✓  |    | ✓       |   |    |    | 3.51      |
|         | ✓       | ✓  |    |         | ✓ |    |    | 3.69      |
|         | ✓       | ✓  | ✓  |         | ✓ |    |    | 3.66      |
|         | ✓       | ✓  | ✓  |         |   | ✓  |    | 3.35      |
|         | ✓       | ✓  | ✓  |         |   |    | ✓  | 3.24      |

Table 3: The test error of ResNeXt-29 (8x64d) with different settings on CIFAR-10. Here, ‘ep.’ indicates epochs. ‘CA’ and ‘CD’ are cosine annealing and curriculum dropout. ‘D’, ‘SD’, and ‘SS’ indicate standard dropout, SpatialDropout, and SelectScale respectively.

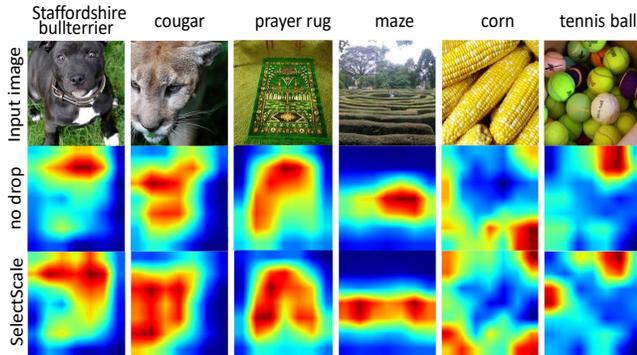


Figure 5: The class activation maps for ResNet-50 on the validation set of ImageNet. These maps are generated by the  $7 \times 7$  feature maps in group 4.

| method            | top-1(%) | top-5(%) |
|-------------------|----------|----------|
| without data drop | 23.89    | 7.06     |
| standard dropout  | 23.97    | 7.12     |
| SpatialDropout    | 23.76    | 7.04     |
| SelectScale       | 23.38    | 6.78     |

Table 4: The test error of ResNet-50 with data drop regularization on ImageNet.

### 3.6 The Patterns Learned by SelectScale

We generate the class activation mapping (CAM) with the features before the global average pooling layer for original ResNet-50 and ResNet-50 with SelectScale. The shown examples in Figure 5 are selected from the validation set. As can be seen, SelectScale can find more patterns. During training, SelectScale continuously changes the learned patterns, which encourages the network to learn new patterns and prevent them from overfitting learned patterns. That is why networks using SelectScale are more robust and perform better.

## 4 Discussion

There are two hyperparameters that need to be tuned in SelectScale. The data retaining rate in SelectScale behaves like that in standard dropout. The network performance will increase as the retaining rate decreases before achieving the best retaining rate. After that, the performance gain will begin to decrease. The other hyperparameter is top rate. The bigger the top rate is, the more features will be changed. In our experiments, most tuned top rates fall in  $\{0.1, 0.2\}$ . We suggest using these two rates if one does not want to tune the top rate.

## 5 Conclusions

In this paper, we present SelectScale to change the important learned features in networks effectively. By adjusting the learned patterns, the networks are forced to learn more patterns from images. Thus the networks with SelectScale become more robust and accurate. Compared to standard dropout and SpatialDropout which drop all features randomly, SelectScale is more effective and introduces less noise.

The method in this paper is fast and parameter-free. It can provide stable improvement for different networks on different datasets. Additionally, our feature selection method does not rely on data drop methods. Given a new data drop method, it is easy to construct a “selecting version” according to our feature ranking strategy. We hope that SelectScale may become a new tool in helping researchers train deep neural networks.

## Acknowledgements

This work is supported by the National Key R&D Program of China under Grant No. 2017YFB1301100, the National Natural Science Foundation of China under Grant Nos. 61976012 and 61772060, and CERNET Innovation Project under Grant No. NGII20170315.

## References

- [Abadi *et al.*, 2016] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *OSDI*, volume 16, pages 265–283, 2016.
- [DeVries and Taylor, 2017] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- [Ghiasi *et al.*, 2018] Golnaz Ghiasi, Tsung-Yi Lin, and Quoc V Le. Dropblock: A regularization method for convolutional networks. *arXiv preprint arXiv:1810.12890*, 2018.
- [Girshick *et al.*, 2014] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014.
- [Graves *et al.*, 2013] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *Acoustics, speech and signal processing (icassp), 2013 IEEE international conference on*, pages 6645–6649. IEEE, 2013.
- [He *et al.*, 2016a] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [He *et al.*, 2016b] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, pages 630–645. Springer, 2016.
- [Huang *et al.*, 2016a] Gao Huang, Zhuang Liu, Kilian Q Weinberger, and Laurens van der Maaten. Densely connected convolutional networks. *arXiv preprint arXiv:1608.06993*, 2016.
- [Huang *et al.*, 2016b] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *European Conference on Computer Vision*, pages 646–661. Springer, 2016.
- [Ioffe and Szegedy, 2015] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [Krizhevsky *et al.*, 2012] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [Long *et al.*, 2015] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [Loshchilov and Hutter, 2016] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- [Morerio *et al.*, 2017] Pietro Morerio, Jacopo Cavazza, Riccardo Volpi, René Vidal, and Vittorio Murino. Curriculum dropout. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3544–3552, 2017.
- [Ren *et al.*, 2015] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [Russakovsky *et al.*, 2015] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [Simonyan and Zisserman, 2014] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [Srivastava *et al.*, 2014] Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [Szegedy *et al.*, 2015] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2015.
- [Szegedy *et al.*, 2016] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2818–2826, 2016.
- [Tompson *et al.*, 2015] Jonathan Tompson, Ross Goroshin, Arjun Jain, Yann LeCun, and Christoph Bregler. Efficient object localization using convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 648–656, 2015.
- [Xie *et al.*, 2016] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. *arXiv preprint arXiv:1611.05431*, 2016.
- [Zagoruyko and Komodakis, 2016] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- [Zhou *et al.*, 2016] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2921–2929, 2016.