

You Look Twice: GaterNet for Dynamic Filter Selection in CNNs

Zhourong Chen^{1,2*}, Yang Li¹, Samy Bengio¹, Si Si¹
¹Google Research, Mountain View

²The Hong Kong University of Science and Technology, Hong Kong
 zchenbb@cse.ust.hk, {liyang, bengio, sisidaisy}@google.com

Abstract

The concept of conditional computation for deep nets has been proposed previously to improve model performance by selectively using only parts of the model conditioned on the sample it is processing. In this paper, we investigate input-dependent dynamic filter selection in deep convolutional neural networks (CNNs). The problem is interesting because the idea of forcing different parts of the model to learn from different types of samples may help us acquire better filters in CNNs, improve the model generalization performance and potentially increase the interpretability of model behavior. We propose a novel yet simple framework called *GaterNet*, which involves a backbone and a gater network. The backbone network is a regular CNN that performs the major computation needed for making a prediction, while a global gater network is introduced to generate binary gates for selectively activating filters in the backbone network based on each input. Extensive experiments on CIFAR and ImageNet datasets show that our models consistently outperform the original models with a large margin. On CIFAR-10, our model also improves upon state-of-the-art results.

1. Introduction

It is widely recognized in neural science that distinct parts of the brain are highly specialized for different types of tasks [20]. It results in not only the high efficiency in handling a response but also the surprising effectiveness of the brain in learning new events. In machine learning, *conditional computation* [3] has been proposed to have a similar mechanism in deep learning models. For each specific sample, the basic idea of conditional computation is to only involve a small portion of the model in prediction. It also means that only a small fraction of parameters needs to be updated at each back-propagation step, which is desirable for training a large model.

*Student of The Hong Kong University of Science and Technology. Work is done when interning at Google.




Figure 1. Model Architecture (better viewed in color). The gater extracts features and generates sparse binary gates for selecting filters in the backbone network in an input-dependent manner.

One line of work that is closely related to conditional computation is *Mixture of Experts* (MoE) [15], where multiple sub-networks are combined via an ensemble using weights determined by a gating module. Particularly, several recent works [25, 28] propose to ensemble a small subset of dynamically selected experts in the model for each input. By doing so, these models are able to reduce computation cost while achieving similar or even better results than baseline models. Note that both the expert architectures and the number of experts in these works are predefined and fixed. Another line of works that resemble conditional computation focus on *dynamic network configuration* [29, 1, 7, 2, 4]. There are no explicitly defined experts in these methods. Rather, they dynamically select the units, layers, or other components in the main model for each input. In these works, one small sub-module is usually added to each position to be configured in the model. That is, each sub-module added is making decisions locally specific to the components it is configuring.

In this paper, we propose a novel framework called *GaterNet* for input-dependent dynamic filter selection in convolutional neural networks (CNNs), as shown in Figure 1. We introduce a dedicated sub-network called *gater*,

which extracts features from input and generates the binary gates needed for controlling filters all at once based on the features. The gating vector is then used to select the filters in the *backbone*¹ network (the main model in our framework), and only the selected filters in the backbone network participate in the prediction and learning. We used a discretization technique called *Improved SemHash* [17] to enable differentiable training of input-dependent binary gates such that the backbone and the gater network can be trained jointly via back-propagation.

Compared to previous works on dynamic network configuration, we use a dedicated sub-network (the gater) for making global decisions on which filters in the backbone network should be used. The decision on each gate (for each filter) is made based on a shared global view of the current input. We argue that such a global gating unit can make more holistic decisions about how to optimally use the filters in the network than local configuration employed by previous work. Note that in [28], a module of the network is used to generate all the gates, which at a glance is similar to our gater. However, there are two important differences. Firstly, [28] is not based on an end-to-end approach. It requires a pre-processing step to cluster classes of samples and assign each cluster to a sub-branch of the network to handle. The assignments provides explicit supervision for training the gating module. Secondly, as mentioned above, the sub-branch architectures and the number of branches are both manually defined and fixed throughout training in [28]. In contrast, in our framework, each sample uses a dynamically determined sub-branch depending on the filters being selected. As a result, our method potentially allows a combinatorial number of choices of sub-branches or experts given the number of filters to be controlled, which is more amenable for capturing complex distribution manifested in the data.

Our experiments on CIFAR [21] and ImageNet [24] classification datasets show that the gater in GaterNet is able to learn effective gating strategies for selecting proper filters. It consistently improves the original model with a significant margin. On CIFAR-10, our method gives better classification results than state-of-the-art models with only 1.2% additional parameters. Our contributions are summarized as follows:

- We propose a new framework for dynamic filter selection in CNNs. The core of the idea is to introduce a dedicated gater network to take a glimpse of the input, and then generate input-dependent binary gates to select filters in the backbone network for processing the input. By using Improved SemHash, the gater network can be trained jointly with the backbone in an end-to-end fashion through back-propagation.
- We conduct extensive experiments on GaterNet, which

show that it consistently improves the generalization performance of deep CNNs without significantly increasing the model complexity. In particular, our models achieve better results than several state-of-the-art models on the CIFAR-10 dataset by only introducing a small fraction of parameters.

- We perform an in-depth analysis about the model behavior for GaterNet, which reveals that GaterNet learns effective gating strategies by being relatively deterministic on the choice of filters to use in shallow layers but using more input-dependent filters in the deep layers.

2. Related Work

The concept of conditional computation is first discussed by Bengio in [3]. Early works on conditional computation focus on how to select model components on the fly. Bengio et al. have studied four approaches for learning stochastic neurons in fully-connected neural networks for conditional selection in [4]. On the other hand, Davis and Arel have used low-rank approximations to predict the sparse activations of neurons at each layer [6]. Bengio et al. have also tested reinforcement learning to optimize conditional computation policies [2].

More recently, Shazeer et al. have investigated the combination of conditional computation with Mixture of Experts on language modeling and machine translation tasks [25]. At each time step in the sequence model, they dynamically select a small subset of experts to process the input. Their models significantly outperformed state-of-the-art models with a low computation cost. In the same vein, Mullapudi et al. have proposed HydraNets that uses multiple branches of networks for extracting features [28]. In this work, a gating module is introduced to generate decisions on selecting branches for each specific input. This method requires a pre-processing step of clustering the ground-truth classes to force each branch to learn features for a specific cluster of classes as discussed in the introduction.

Dynamic network configuration is another type of conditional computation that has been studied previously. In this line of works, no parallel experts are explicitly defined. Instead, they dynamically configure a single network by selectively activating model components such as units and layers for each input. Adaptive Dropout is proposed by Ba and Frey to dynamically learn a dropout rate for each unit and each input [1]. Denoyer and Ludovic have proposed a tree structure neural network called Deep Sequential Neural Network [7]. A path from the root to a leaf node in the tree represents a computation sequence for the input, which is also dynamically determined for each input. Recently, Veit and Belongie [29] have proposed to skip layers in ResNet [10] in an input-dependent manner. The resulting model is performing better and also more robust to adver-

¹The term *backbone* is also used in object detection and TSE-Net [5].

serial attack than the original ResNet, which also leads to reduced computation cost.

Previous works have also investigated methods that dynamically re-scale or calibrate the different components in a model. The fundamental difference between these methods and dynamic network configuration is that they generate a real-valued vector for each input, instead of a binary gate vector for selecting network components. SE-Net proposed by Hu et al. [12] re-scales the channels in feature maps on the fly and achieves state-of-the-art results on ImageNet classification dataset. Stollenga et al. [26] have also proposed to go through the main model for multiple passes. The features resulting from each pass (except the last) are used to generate a real-valued vector for re-scaling the channels in the next pass. In contrast to these works, our gater network generates binary decisions to dynamically turn on or off filters depending on each input.

3. GaterNet

Our model contains two convolutional neural sub-networks, namely the *backbone network* and the *gater network* as illustrated in Figure 1. Given an input, the gater network decides the set of filters in the backbone network for use while the backbone network does the actual prediction. The two sub-networks are trained in an end-to-end manner via back-propagation.

3.1. Backbone

The backbone network is the main module of our model, which extracts features from input and makes the final prediction. Any existing CNN architectures such as ResNet [10], Inception [27] and DenseNet [13] can be readily used as the backbone network in our GaterNet.

Let us first consider a standalone backbone CNN without the gater network. Given an input image x , the output of the l -th convolutional layer is a 3-D feature map $O^l(x)$. In a conventional CNN, $O^l(x)$ is computed as:

$$O_i^l(x) = \phi(F_i^l * I^l(x)), \quad (1)$$

where $O_i^l(x)$ is the i -th channel of feature map $O^l(x)$, F_i^l is the i -th 3-D filter, $I^l(x)$ is the 3-D input feature map to the l -th layer, ϕ denotes the element-wise nonlinear activation function, and $*$ denotes convolution. In general cases without the gater network, all the filters F_i^l in the current layer are applied to $I^l(x)$, resulting in a *dense* feature map $O^l(x)$. The loss for training such a CNN for classification is $L = -\log P(y|x, \theta)$ for a single input image, where y is the ground-truth label and θ denotes the model parameters.

3.2. Gater

In contrast to the backbone, the gater network is an assistant of the backbone and does not learn any features directly

used in the prediction. Instead, the gater network processes the input to generate an input-dependent *gating mask*—a *binary vector*. The vector is then used to dynamically select a particular subset of filters in the backbone network for the current input. Specifically, the gater network learns a function as below:

$$G(x) = D(E(x)). \quad (2)$$

Here, E is an image feature extractor defined as $E : x \rightarrow f, x \in \mathbb{R}^{h' \times w' \times c'}, f \in \mathbb{R}^h$, with h', w', c' being the height, width and channel number of an input image respectively, and h being the number of features extracted. D is a function defined as $D : f \rightarrow g, f \in \mathbb{R}^h, g \in \{0, 1\}^c$, where c is the total number of filters in the backbone network. More details about function E and D will be discussed in Section 3.2.1 and Section 3.2.2 respectively.

From the above definition we can see that, the gater network learns a function which maps input x to a binary gating vector g . With the help of g , we reformulate the computation of feature map $O^l(x)$ in Equation (1) as below:

$$O_i^l(x) = \begin{cases} \mathbf{0}, & \text{if } g_i^l = 0 \\ \phi(F_i^l * I^l(x)), & \text{if } g_i^l = 1 \end{cases} \quad (3)$$

Here g_i^l is the entry in g corresponding to the i -th filter at layer l , and $\mathbf{0}$ is a 2-D feature map with all its elements being 0. That is, the i -th filter will be applied to $I^l(x)$ to extract features only when $g_i^l = 1$. If $g_i^l = 0$, the i -th filter is skipped and $\mathbf{0}$ is used as the output instead. When g^l is a sparse binary vector, a large subset of filters will be skipped, resulting in a sparse feature map. In this paper, we implement the computation in Equation (3) by masking the output channels using the binary gates:

$$O_i^l(x) = \phi(F_i^l * I^l(x)) \cdot g_i^l \quad (4)$$

In the following subsections, we will introduce how we design the functions E and D in Equation (2) and how we enable end-to-end training through the binary gates.

3.2.1 Feature Extractor

Essentially, the function $E(x)$ in Equation (2) is a feature extractor which takes an image x as input and outputs a feature vector f . Similar to the backbone network, any existing CNN architectures can be used here to learn the function $E(x)$. There are two main differences compared with the backbone network: (1) The output layer of the CNN architecture is removed such that it outputs features for use in the next step. (2) A gater CNN does not necessarily need to be as complicated as the one for the backbone. One reason is that the gater CNN is supposed to obtain a brief view of the input. Having an over-complicated gater network may encounter various difficulties in computation cost and optimization. Another reason is to avoid the gater network accidentally taking over the task that is intended for the backbone network.

3.2.2 Features to Binary Gates

Fully-Connected Layers with Bottleneck As defined in Equation (2), the function $D(f)$ needs to map the vector f of size h to a binary vector g of size c . We first consider using fully-connected layers to map f to a *real-valued* vector g' of size c . If we use one single layer to project the vector, the projection matrix would be of size $h \times c$. This can be very large when h is thousands and c is tens of thousands. To reduce the number of parameters in this projection, we use two fully-connected layers to fulfill the projection. The first layer projects f to a bottleneck of size b , followed by the second layer mapping the bottleneck to g' . In this way, the total number of parameters becomes $(h+c) \times b$. This can be significantly smaller than $h \times c$ when b is much smaller than h and c . We ignore bias parameters here for simplicity.

In summary, the real-valued vector g' is computed as:

$$\begin{aligned} f' &= \text{ReLU}(\text{BatchNorm}(FC_1(f))) \\ g' &= FC_2(f') \end{aligned}$$

where FC_1 and FC_2 denotes the two linear projections, ReLU denotes the non-linear activation function in [23], and BatchNorm means batch normalization [14].

Improved SemHash So far, one important question still remains unanswered: how to generate binary gates g from g' such that we can back-propagate the error through the discrete gates to the gater? In this paper, we adopt a method called *Improved SemHash* [17, 18].

During training, we first draw noise from a c -dimensional Gaussian distribution with mean 0 and standard deviation 1. The noise ϵ is added to g' to get a noisy version of the vector: $g'_\epsilon = g' + \epsilon$. Two vectors are then computed from g'_ϵ :

$$g_\alpha = \sigma'(g'_\epsilon) \text{ and } g_\beta = \mathbf{1}(g'_\epsilon > 0)$$

where σ' is the saturating sigmoid function [19, 16]:

$$\sigma'(x) = \max(0, \min(1, 1.2\sigma(x) - 0.1))$$

with σ being the sigmoid function. Here, g_α is a real-valued gate vector with all the entries falling in the interval $[0.0, 1.0]$, while g_β is a binary vector. We can see that, g_β has the desirable binary property that we want to use in our model, but the gradient of g_β w.r.t g'_ϵ is zero for most values of g'_ϵ . On the other hand, the gradient of g_α w.r.t g'_ϵ is well defined, but g_α is not a binary vector. In forward propagation, we randomly use $g = g_\alpha$ for half of the training samples and use $g = g_\beta$ for the rest of the samples. When g_β is used, we follow the solution in [17, 18] and define the gradient of g_β w.r.t g'_ϵ to be the same as the gradient of g_α w.r.t g'_ϵ in the backward propagation.

The above procedure is designed for the sake of easy training. Evaluation and inference are different to the training phase in two aspects. Firstly, we skip the step of drawing noise and always set $\epsilon = 0$. Secondly, we always use

the discrete gates $g = g_\beta$ in forward propagation. That is, the gate vector is always binarized in evaluation and inference phase. The interested readers are referred to [17, 18] for more intuition behind Improved SemHash.

We use binary gates other than attention [30] or other real-valued gates for two reasons. Firstly, binary gates can completely deactivate some filters for each input, and hence those filters will not be influenced by the irrelevant inputs. This may lead to training better filters than real-valued gates. Secondly, discrete gates open the opportunity for model compression in the future.

Sparse Gates To encourage the gates g to be sparse, we introduce a L_1 regularization term into the training loss:

$$L = -\log P(y|x, \theta) + \lambda \frac{\|G(x)\|_1}{c}$$

where λ is the weight for the regularization term and c is the size of g . Note that the backbone network receives no gradients from the second term, while the gater network receives gradients from both the two terms.

3.3. Pre-training

While our model architecture is straightforward, there are several empirical challenges to train it well. First, it is difficult to learn these gates, which are discrete latent representations. Although Improved SemHash has been shown to work well in several previous works, it is unclear whether the approximation of gradients mentioned above is a good solution in our model. Second, the introduction of gater network into the model has essentially changed the optimization space. The current parameter initialization and optimization technique may not be suitable for our model. We leave the exploration of better binarization, initialization and optimization techniques to our future works. In this paper, we always initialize our backbone network and gater network from networks pre-trained on the same task, and empirically find it works well with a range of models.

4. Experiments

We first conduct preliminary experiments on CIFAR [21] with ResNet [10, 11], which gives us a good understanding about the performance improvements our method can achieve and also the gating strategies that our gater is learning. Then we apply our method to state-of-the-art models on CIFAR-10 and show that we consistently outperform these models. Lastly, we move on to a large-scale classification dataset, ImageNet 2012 [24], and show that our method significantly improves the performance of large models, such as ResNet and Inception-v4 [27], as well.

4.1. Datasets

CIFAR-10 and CIFAR-100 contain natural images belonging to 10 and 100 classes respectively. There are 50,000

training and 10,000 test images. We randomly hold out 5,000 training images as a validation set. All the final results reported on test images are using models trained on the complete training set. The raw images are with 32×32 pixels and we normalize them using the channel means and standard deviations. Standard data augmentation by random cropping and mirroring are applied to the training set. ImageNet 2012 classification dataset contains 1.28 million training images and 50,000 validation images of 1,000 classes. We use the same data augmentation method as the original papers of the baseline models in Table 3. The images are of 224×224 and 299×299 in ResNet and Inception-v4 respectively.

4.2. Cifar-10 and CIFAR-100

4.2.1 Preliminary Experiments with ResNet

We first validate the effectiveness of our method using ResNet as the backbone network on CIFAR-10 and CIFAR-100 datasets. We consider a shallow version, ResNet-20, and two deep versions, ResNet-56 and ResNet-164² to gain a better understanding on how our gating strategy can help models with varying capacities. All our gated models employ ResNet-20 as the gater network. Table 1 shows the comparison with baseline models on the test set. *ResNet-Wider* is the ResNet with additional filters at each layer such that it contains roughly the same number of parameters as our model. *ResNet-SE* is the ResNet with squeeze-and-excitation block [12]. The *Gated Filters* column shows the number of filters under consideration in our models.

Classification Results From the table we can see that, our model consistently outperforms the original ResNet with a significant margin. On CIFAR100, the error rate of ResNet-164 is reduced by 1.83%.

It is also evident that, our model is performing better than ResNet-SE in all cases. Note that our gater network is generating *binary gates* for the backbone network channels, while ResNet-SE is re-scaling the channels. It is interesting that, although our method is causing more information loss in the forward pass of backbone network due to the sparse discrete gates, our model still achieves better generalization performance than ResNet and ResNet-SE. This to some extent validates our assumption that only a subset of filters are needed for the backbone to process an input sample.

Ablation Analysis on Model Size In all cases, ResNet-Wider is better than the original ResNet as well. ResNet-20-Wider is even the best among all the shallow models. We hypothesize that ResNet-20 is suffering from underfitting due to its small amount of filters and hence adding additional filters significantly improves the model. On the other hand, although ResNet-20-Gated has a similar number of parameters as ResNet-20-Wider, a significant portion

(about a half) of its parameters belongs to the gater network, rather than directly participating in prediction, and ResNet-20-Gated still performed on par with ResNet-20-Wider.

The backbone network in ResNet-20-Gated suffers from underfitting due to the lack of effective filters. The comparison among the deep models validates our hypothesis. ResNet-50 and ResNet-164 contain many more filters than ResNet-20, and adding filters to them shows only a minor improvement (see ResNet-50-Wider and ResNet-164-Wider). In these cases, our models show a significant improvement over the wider models and are the best among all the deep models on both datasets. The comparison with ResNet-Wider shows that the effectiveness of our model is not solely due to the increase of parameter number, but mainly due to our new gating mechanism.

Complexity It appears to be an issue at a glance if a comprehensive gater network is needed to assist a backbone network, as it may greatly increase the number of parameters. However, our experiments show that the gater network does not need to be complex, and as a matter of fact, it can be much smaller than the backbone network (see Table 1). Although the number of filters (in the backbone network) under consideration varies from 336 to 7200, the results show that a simple gater network such as ResNet-20 is powerful enough to learn input-dependent gates for the three models that have a wide range of model capacity. As such, when the backbone network is large (where our method shows more significant improvements over baselines), the parameter overhead introduced by the gater network becomes small. For example, ResNet-164-Gated has only 20% more parameters than ResNet-164. In contrast, in other more complicated backbone networks such as DenseNet and Shake-Shake, this overhead is reduced to 1.2% as shown in Table 2. Consequently, the complexity and the number of additional parameters that our method brings to an existing model is relatively small, especially to large models.

Gate Distribution One question that would naturally occur is how the distribution of the learned gates looks like. Firstly, it is possible that the gater network is just randomly pruning the backbone network and introducing regularization effects similar to dropout into the backbone. It is also possible that the gates are always the same for different samples. Secondly, the generated gates may give us good insights into the importance of filters at different layers.

To answer these questions, we analyzed the gates generated by the gater network in ResNet-164-Gated. We first conduct forward propagation in the gater network on CIFAR-10 test set and collect gates for all the test samples. As expected, three types of gates emerge: gates that are always on for all the samples, gates that are always off, and gates that can be on or off conditioned on the input, i.e., input-dependent gates. We show the percentage of the three types of gates at different depth in Figure 2. We can see

²Our ResNet-164 is slightly different to the one in [11]. The number of filters in the first group of residual units are 16, 4, 16 respectively.

Table 1. Classification error rates on the CIFAR-10 and CIFAR-100 test set. All the methods are with data augmentation. ResNet-Wider is the ResNet with additional filters at each layer such that it contains roughly the same number of parameters as our model. ResNet-SE is the ResNet with squeeze-and-excitation blocks. All the ResNet-Gated models are using ResNet-20 as the gater network. The Gated Filters column shows the number of filters subject to gating in our models. All the baseline results are from our reimplementation.

	Gated Filters	Cifar10		Cifar100	
		Param	Error Rates %	Param	Error Rates %
ResNet-20 [10]	-	0.27M	8.06	0.28M	32.39
ResNet-20-Wider	-	0.56M	6.85	0.57M	30.08
ResNet-20-SE [12]	-	0.28M	7.81	0.29M	31.22
ResNet-20-Gated (Ours)	336	0.55M	6.88 (↓1.18)	0.60M	30.79 (↓1.60)
ResNet-56 [10]	-	0.86M	6.74	0.86M	28.87
ResNet-56-Wider	-	1.08M	6.72	1.09M	28.39
ResNet-56-SE [12]	-	0.88M	6.27	0.89M	28.00
ResNet-56-Gated (Ours)	1,008	1.14M	5.72 (↓1.02)	1.14M	27.71 (↓1.16)
ResNet-164 [11]	-	1.62M	5.61	1.64M	25.39
ResNet-164-Wider	-	2.04M	5.57	2.07M	24.80
ResNet-164-SE [12]	-	2.00M	5.51	2.02M	23.83
ResNet-164-Gated (Ours)	7,200	1.96M	4.80 (↓0.81)	1.98M	23.56 (↓1.83)




Figure 2. The distribution of Gates in each layer for ResNet-164-Gated on Cifar-10 Test Set. There are totally 54 residual units.

that, a large subset (up to 68.75%) of the gates are always off at the shallow residual blocks. As the backbone network goes deeper, the proportion of always-on and input-dependent gates increases gradually. In the last two residual blocks, input-dependent gates become the largest subset of gates with percentages of around 45%. The phenomenon is consistent with the common belief that shallow layers are usually extracting low-level features which are essential for all kinds of samples, while deep layers are extracting high-level features which are very sample-specific.

Although the above figures show that the gater network is learning input-dependent gates, it does not show how often that these gates are on/off. For example, a gate that is on for only one test sample but off for the rest would also appear input-dependent. To investigate this further, we collect all the input-dependent gates and plot the distribution of number of times that they are on in Figure 3. There are totally 1567 input-dependent gates out of the total number of 7200 gates for the backbone network. While many of these gates remain in one state—either on or off—in most of the time, there are 1,124 gates that switch on and off more frequently—they are activated for 100 ~ 9900 samples out




Figure 3. Distribution: X-axis is the number of times an input-dependent gate is on, while Y-axis is the number of gates.




Figure 4. Distribution: X-axis is the number of gates on, while Y-axis is the number of samples.

of the 10,000 test samples.

We also examined how many gates are fired when processing each test example. The maximum and minimum number of fired gates per sample is 5380 and 5506 respectively. The average number is around 5453. The number of gates used each time seems to obey a normal distribution (see Figure 4).

Lastly, we want to investigate what gating strategy has been learned by the gater network. To do so, we represent the filter usage of each test sample as a 7200-dimensional binary vector where each element in the vector represents if the corresponding gate is on (1) or off (0). We collect the filter usage vector of each sample and reduce the dimension




Figure 5. Visualization of the high dimensional gate vectors. PCA and t-SNE are applied to reduce the vector dimension to 2. Dots with the same color corresponds to test samples with the same label. Better viewed in color.

of these vectors from 7200 to 400 using Principal Component Analysis (PCA). We then project these vectors onto a 2-dimensional space via t-SNE [22] (see Figure 5). Interestingly, we find samples of the same class tend to use similar gates. In the figure, each color of dots represents a ground-truth label. This shows that the gater network learned to turn on similar gates for samples from the same class—hence similar parts of the backbone network are used to process the samples from the class. On the other hand, we found the clusters in Figure 5 is still far from perfectly setting samples from different labels apart. It is indeed a good evidence that the gater network doesn’t accidentally take over the prediction task that the backbone network is intended to do, which is what we want to avoid. We want the gater network to focus on learning to make good decisions on which filters in the backbone network should be used. From this analysis, we can see that the experiments are turned out as we expected and the backbone network still does the essential part of prediction for achieving the high accuracy.

We can draw the following conclusions from the above observations and analyses:

- The gater network is capable of learning effective gates for different samples. It tends to generate similar gates for samples from the same class (label).
- The residual blocks at shallow layers are more redundant than those at deep layers.
- Input-dependent features are more needed at deep layers than at shallow layers.

4.2.2 State-of-the-Art on CIFAR-10

Next we test the performance of our method with state-of-the-art models, Shake-Shake [9] and DenseNet [13], on CIFAR-10 dataset. We use Shake-Shake and DenseNet as the backbone network and ResNet-20 again as the gater

network to form our models respectively. Table 2 summarizes the comparison of our models with the original models. The gater network in our method consistently improves the state-of-the-art backbone network without significantly increasing the number of parameters. One of our models, *Shake-Shake-Gated 26 2x96d*, has only 1.2% more parameters than the corresponding baseline model. Another interesting finding is that, with the assistance of the gater network, *DenseNet-BC-Gated ($L = 250, K = 24$)* is even performing better than both *DenseNet-BC ($L = 190, K = 40$)* and *DenseNet-BC-Gated ($L = 190, K = 40$)*, although it has much fewer parameters.

Note that in [8], it is shown when *Shake-Shake 26 2x96d* is combined with a data pre-processing technique called *cutout*, it can achieve 2.56% error rate on CIFAR-10 test set. The technique is orthogonal to our method and can also be combined with our method to give better results.

4.3. ImageNet

To test the performance of our method on large datasets, we apply our method to models for ImageNet. We use ResNet [10] and Inception-v4 [27] as the backbone network and ResNet-18 [10] as the gater network to form our models. Table 3 shows the classification results on ImageNet validation set with baselines similar to the settings in Table 1. We can see that, our method improves all the models by 0.52% ~ 1.85% in terms of top-1 error rate, and 0.14% ~ 0.78% in terms of top-5 error rate. Note that [29] proposes to dynamically skip layers in ResNet-101, and the top-1 and top-5 error rates of their model are 22.63% and 6.26% respectively. Our ResNet-101-Gated achieves 21.51% and 5.72% on the same task, which is apparently much better than their model. In addition, there are also two interesting findings:

- The performance of ResNet-101 is significantly boosted with the help of the gater network. ResNet-101-Gated is even performing better than ResNet-152 using much fewer layers.
- Similar to the results on CIFAR datasets, ResNet-Wider is performing well when the original model is shallow and small, but is outperformed by our models when the original model contains enough filters.

4.4. Implementation Details

We train the baseline models by following the training schemes in the original papers. We pre-train the backbone and the gater network on the target task separately to properly initialize the weights. The training scheme here includes training configurations such as number of training epochs, learning rate, batch size, weight decay and so on.

After pre-training, we train the backbone and the gater network jointly as a single model. In addition to following the original training scheme for each backbone archi-

Table 2. Classification error rates on the CIFAR-10 test set. All the methods use data augmentation during training. All our models are using ResNet-20 as the gater network. The Gated Filters column shows the number of filters subject to gating in our models. All the baseline results are from our reimplementation.

	Gated Filters	Param	Error Rates %
DenseNet-BC ($L = 100, k = 12$) [13]	-	0.77M	4.48
DenseNet-BC-Gated ($L = 100, k = 12$, Ours)	540	1.05M	4.03
DenseNet-BC ($L = 250, k = 24$) [13]	-	15.32M	3.61
DenseNet-BC-Gated ($L = 250, k = 24$, Ours)	2,880	15.62M	3.31
DenseNet-BC ($L = 190, k = 40$) [13]	-	25.62M	3.52
DenseNet-BC-Gated ($L = 190, k = 40$, Ours)	3,600	25.93M	3.39
Shake-Shake 26 2x64d [9]	-	11.71M	3.05
Shake-Shake-Gated 26 2x64d (Ours)	3,584	12.01M	2.89
Shake-Shake 26 2x96d [9]	-	26.33M	2.82
Shake-Shake-Gated 26 2x96d (Ours)	5,376	26.65M	2.64

Table 3. Single-crop error rates on the ImageNet 2012 validation set. ResNet-Wider is the ResNet with additional filters at each layer such that it contains roughly the same number of parameters as our model. ResNet-SE is the ResNet with squeeze-and-excitation units. All the ResNet-Gated models are using ResNet-18 as the gater network. The Gated Filters column shows the number of filters subject to gating in our models. All the baseline results are from our reimplementation, except [†] from the original paper.

	Gated Filters	Parameters	Top-1 Error %	Top-5 Error %
ResNet-34 [10]	-	21.80M	26.56	8.48
ResNet-34-Wider	-	33.89M	25.36	7.91
ResNet-34-SE [12]	-	21.96M	26.08	8.30
ResNet-34-Gated (Ours)	3,776	34.08M	26.04 (\downarrow 0.52)	8.34 (\downarrow 0.14)
ResNet-101 [10]	-	44.55M	23.36	6.56
ResNet-101-Wider	-	59.17M	21.89	6.05
ResNet-101-SE [12]	-	49.33M	22.38 [†]	6.07 [†]
ResNet-101-Gated (Ours)	32,512	64.21M	21.51 (\downarrow 1.85)	5.78 (\downarrow 0.78)
ResNet-152 [10]	-	60.19M	22.34	6.22
ResNet-152-Wider	-	81.37M	21.50	5.67
ResNet-152-SE [12]	-	66.82M	21.57 [†]	5.73 [†]
ResNet-152-Gated (Ours)	47,872	83.80M	21.19 (\downarrow 1.15)	5.45 (\downarrow 0.77)
Inception-v4 [27]	-	44.50M	20.33	4.99
Inception-v4-Gated (Ours)	16,608	61.67M	19.64 (\downarrow 0.69)	4.80 (\downarrow 0.19)

ture, we introduce a few minor modifications. Firstly, we increase the number of training epochs for DenseNet-Gated and Shake-Shake-Gated by 20 and 30 respectively as they seem to converge slowly at the end of training. Secondly, we set the initial learning rate for DenseNet-Gated and Shake-Shake-Gated to a smaller value, 0.05.

Note that not all the filters in a backbone network are subject to gating in our experiments. When ResNet is used as the backbone, we apply filter selection to the last convolutional layer in each residual unit, which is similar to SE-block [12]. As for DenseNet, we apply filter selection to all the convolutional layers except the first in each dense block. There are multiple residual branches in each residual block in Shake-Shake. We apply filter selection to the last convolutional layer in each branch. In Inception-v4, there are many channel concatenation operations. We apply filter selection to feature maps after the concatenation operations.

For all our models on CIFAR, we set the size of the bottleneck layer to 8. ResNet-34-Gated, ResNet-101-Gated

and Inception-v4-Gated use a bottleneck size of 256, while ResNet-152-Gated uses 1024.

5. Conclusions

In this paper, we have proposed GaterNet, a novel architecture for input-dependent filter selection in CNNs. It involves two distinct components: a backbone network that conducts actual prediction and a gater network that decides which part of backbone network should be used for processing each input. Extensive experiments on CIFAR and ImageNet indicate that our models consistently outperform the original models with a large margin. On CIFAR-10, our model improves upon state-of-the-art results. We have also performed an in-depth analysis about the model behavior that reveals an intuitive gating strategy learned by the gater network.

6. Acknowledgments

We want to thank Sergey Ioffe, Noam Shazeer, Zhenhai Zhu and anonymous reviewers for their insightful feedback.

References

- [1] Jimmy Ba and Brendan Frey. Adaptive dropout for training deep neural networks. In *NIPS*, pages 3084–3092, 2013.
- [2] Emmanuel Bengio, Pierre-Luc Bacon, Joelle Pineau, and Doina Precup. Conditional computation in neural networks for faster models. *arXiv preprint arXiv:1511.06297*, 2015.
- [3] Yoshua Bengio. Deep learning of representations: Looking forward. In *Statistical Language and Speech Processing*, pages 1–37, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [4] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- [5] Zhoung Chen, Xiaopeng Li, and Nevin L. Zhang. Learning sparse deep feedforward networks via tree skeleton expansion. *CoRR*, abs/1803.06120, 2018.
- [6] Andrew Davis and Itamar Arel. Low-rank approximations for conditional feedforward computation in deep neural networks. *arXiv preprint arXiv:1312.4461*, 2013.
- [7] Ludovic Denoyer and Patrick Gallinari. Deep sequential neural network. In *Deep Learning and Representation Learning Workshop, NIPS 2014*, 2014.
- [8] Terrance Devries and Graham W. Taylor. Improved regularization of convolutional neural networks with cutout. *CoRR*, abs/1708.04552, 2017.
- [9] Xavier Gastaldi. Shake-shake regularization. In *ICLR Workshop*, 2017.
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *ECCV*. Springer, 2016.
- [12] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018.
- [13] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *CVPR*, 2017.
- [14] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.
- [15] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991.
- [16] Łukasz Kaiser and Samy Bengio. Can active memory replace attention? In *NIPS*, pages 3781–3789, 2016.
- [17] Łukasz Kaiser and Samy Bengio. Discrete autoencoders for sequence models. *arXiv preprint arXiv:1801.09797*, 2018.
- [18] Łukasz Kaiser, Samy Bengio, Aurko Roy, Ashish Vaswani, Niki Parmar, Jakob Uszkoreit, and Noam Shazeer. Fast decoding in sequence models using discrete latent variables. In *ICML*, 2018.
- [19] Łukasz Kaiser and Ilya Sutskever. Neural gpus learn algorithms. *arXiv preprint arXiv:1511.08228*, 2015.
- [20] E.R. Kandel, J.H. Schwartz, and T.M. Jessell. *Principles of neural science*. International edition. Elsevier, 1991.
- [21] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, 2009.
- [22] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- [23] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, pages 807–814, 2010.
- [24] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [25] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *ICLR*, 2017.
- [26] Marijn F Stollenga, Jonathan Masci, Faustino Gomez, and Jürgen Schmidhuber. Deep networks with internal selective attention through feedback connections. In *NIPS*, pages 3545–3553, 2014.
- [27] Christian Szegedy, Sergey Ioffe, and Vincent Vanhoucke. Inception-v4, inception-resnet and the impact of residual connections on learning. In *AAAI*, 2017.
- [28] Ravi Teja Mullapudi, William R. Mark, Noam Shazeer, and Kayvon Fatahalian. Hydranets: Specialized dynamic architectures for efficient inference. In *CVPR*, June 2018.
- [29] Andreas Veit and Serge J. Belongie. Convolutional networks with adaptive inference graphs. In *ECCV*, volume 11205 of *Lecture Notes in Computer Science*, pages 3–18. Spribe2013adaptivenger, 2018.
- [30] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, pages 2048–2057, 2015.