

Filter Response Normalization Layer: Eliminating Batch Dependence in the Training of Deep Neural Networks

Saurabh Singh Shankar Krishnan
Google Research

{saurabhsingh, skrishnan}@google.com

Abstract

*Batch Normalization (BN) uses mini-batch statistics to normalize the activations during training, introducing dependence between mini-batch elements. This dependency can hurt the performance if the mini-batch size is too small, or if the elements are correlated. Several alternatives, such as Batch Renormalization and Group Normalization (GN), have been proposed to address this issue. However, they either do not match the performance of BN for large batches, or still exhibit degradation in performance for smaller batches, or introduce artificial constraints on the model architecture. In this paper we propose the Filter Response Normalization (FRN) layer, a novel combination of a normalization and an activation function, that can be used as a replacement for other normalizations and activations. Our method operates on each activation channel of each batch element independently, eliminating the dependency on other batch elements. Our method outperforms BN and other alternatives in a variety of settings **for all batch sizes**. FRN layer performs ≈ 0.7 - 1.0% better than BN on top-1 validation accuracy with large mini-batch sizes for Imagenet classification using InceptionV3 and ResnetV2-50 architectures. Further, it performs $> 1\%$ better than GN on the same problem in the small mini-batch size regime. For object detection problem on COCO dataset, FRN layer outperforms all other methods by at least 0.3 - 0.5% in all batch size regimes.*

1. Introduction

Batch normalization (BN) [18] is a cornerstone of current high performing deep neural network models. One often discussed drawback of BN is its reliance on sufficiently large batch sizes [17, 31, 36]. When trained with small batch sizes, BN exhibits a significant degradation in performance. This issue has been attributed to training and testing discrepancy arising from BN’s reliance on stochastic mini-batches [31]. As a result, several approaches have been proposed that ameliorate the issues due to stochasticity [17, 31] or offer alternatives [2, 36] by removing batch dependence. How-




Figure 1: Our method consistently outperforms other normalization methods, even at the largest batch size where other methods struggle in comparison to Batch Normalization (see inset). The figure reports the validation performance of ResNetV2-50 models trained using 8 GPUs with different batch sizes on ImageNet.

ever, these approaches don’t match the performance of BN for large batch sizes (Figure 1). Further, either they still exhibit a degradation in performance for smaller batch sizes e.g. Batch Renormalization, or introduce constraints on the model architecture and size e.g. Group Normalization requires number of channels in a layer to be multiples of an ideal group size, such as 32. In this work we propose *Filter Response Normalization (FRN) layer*, consisting of a normalization and activation function, that eliminates these shortcomings altogether. Our method does not have any batch dependence, as it operates on each activation channel (filter response) of each batch sample independently, and outperforms BN and alternatives in a wide variety of evaluation settings. For example, in Figure 1, FRN layer outperforms other approaches by more than 1% at *all batch sizes* for ResNetV2-50 on ImageNet classification.

The reliance of BN on large batch sizes hinders the exploration of higher capacity models due to significantly higher memory requirements and imposes limitations on the performance of tasks that need to process larger inputs. For

example, object detection and segmentation perform better with higher resolution inputs; similarly, video data tends to be very high dimensional. As a result, these systems are forced to trade-off between model capacity and ability to train with larger batch sizes. As evidenced in Figure 1, our method addresses this by maintaining a consistent performance across a range of batch sizes.

FRN layer consists of two novel components that work together to yield high performance: 1) *Filter Response Normalization (FRN)*, a normalization method that independently normalizes the responses of each filter for each batch element by dividing them by the square root of their uncentered second moment, *without performing any mean subtraction*, and 2) *Thresholded Linear Unit (TLU)*, a pointwise activation that is parameterized by a learned rectification threshold, allowing for activations that are biased away from zero. FRN layer outperforms BN by more than **0.7-1.0%** with large mini-batch sizes for Imagenet classification using InceptionV3 and ResnetV2-50 architectures. Further, it performs $> 1\%$ better than GN on the same problem in the small mini-batch size regime. For object detection on COCO dataset, FRN layer outperforms all other methods by at least **0.3-0.5%** in all batch size regimes. Lastly, FRN layer maintains a consistent performance across all the batch sizes that we tested. In summary, **the proposed FRN layer does not rely on other batch elements or channels for normalization, yet outperforms BN and other alternatives for all batch sizes and in a variety of settings.**

Contributions:

1. Filter Response Normalization (FRN), a normalization method that enables models trained with per-channel normalization to achieve high accuracy.
2. Thresholded Linear Unit (TLU), an activation function to use with FRN that further improves accuracy and outperforms BN at all batch sizes without any batch dependency. We refer to this combination as FRN layer.
3. Several insights and practical considerations that lead to the success of the combination of FRN and TLU.
4. A detailed experimental study comparing popular normalization methods on large scale image classification and object detection tasks using a variety of architectures.

2. Related work

Normalization of training data has been known to aid in optimization. For example, whitening of inputs is a common practice for training shallow models such as Support Vector Machines and Logistic regression. Similarly, for training deep networks, normalization of inputs and intermediate representations has been recommended for efficient learning [9, 20, 21]. Batch Normalization (BN) [18] accelerates learning and enables training of very deep neural network

architectures by stabilizing the intermediate feature distributions. Stabilization is achieved by normalizing each activation channel independently, using the mean and variance statistics computed for that channel over the entire mini-batch. However, BN exhibits a dramatic degradation in performance when trained with smaller mini-batches [31, 36]. Several approaches have been proposed to address this shortcoming, and can be grouped into two major categories: 1) Methods that reduce the train-test discrepancy in batch normalized models, 2) Sample based normalization methods that avoid batch normalization.

Methods reducing train-test discrepancy in batch normalization. Ioffe [17] notes that the discrepancy between the statistics that are used for normalization during training and testing may arise from the stochasticity due to small mini-batches and bias due to non-iid samples. They propose Batch Renormalization (BR) to reduce this discrepancy by constraining the mini-batch moments to a specific range, limiting the variation in mini-batch statistics during training. A key benefit of this approach is that the test time evaluation scheme of a model trained with BR is exactly the same as that for a model trained with BN. In comparison, EvalNorm [31] does not modify the training scheme. Instead, it proposes a correction to the normalization statistics for use only during evaluation. The major advantage of this method is that the model does not need to be retrained. However, both these methods still exhibit a degradation in performance for small mini-batches. Another approach is to engineer systems that can circumvent the issue by distributing larger batches across GPUs for tasks that require large inputs [25]. However, this approach requires considerable GPU infrastructure.

Methods avoiding normalization using mini-batches. Several approaches sidestep the issues encountered by BN by not relying on the stochastic mini-batch altogether [2, 35, 36]. Instead, the normalization statistics are computed from the sample itself. Layer Normalization (LN) [2] computes the normalization statistics from the entire layer i.e. using all the activation channels. In contrast, like BN, Instance Normalization (IN) [35] computes the normalization statistics for each channel independently, but only from the sample being normalized, as opposed to the entire batch, as BN does. IN was shown to be useful for style transfer applications, but was not successfully applied for recognition. Group Normalization (GN) [36] fills the middle ground between the two. It computes the normalization statistics over groups of channels. The ideal group size is experimentally determined. While, GN doesn't show performance degradation for smaller batch sizes, it performs worse than BN for larger mini-batches (See Figure 1 here and Figure 1 in [36]). In addition, the size of groups required by GN imposes a constraint on the network size and architecture as every normalized layer needs to have number of channels that are a




Figure 2: A schematic of the proposed FRN Layer.

multiple of the ideal group size determined by GN.

Other approaches. Weight Normalization [28] proposes a reparameterization of the filters in terms of a direction and a scale and reports accelerated convergence. Normalization Propagation [1] uses idealized moment estimates to normalize every layer. Refer to Ren et al. [26] for a unifying view of various normalization approaches. Divisive normalization (DN) [6, 16] has been proposed to normalize each activation with a function of the neighboring activations and has been studied in a variety of contexts including density modeling [3], image compression [4, 5], sensitivity maximization [7], distributed neural representation [29] and attention [27], among others.

3. Approach

Our goal is to eliminate the batch dependency in the training of deep neural networks without sacrificing the performance gains of BN at large batch sizes. We start this section with the main details of our proposal and follow that with a discussion of the rationale behind our approach.

3.1. Filter Response Normalization with Thresholded Activation

We will assume for the purpose of exposition that we are dealing with the feed-forward convolutional neural network. We follow the usual convention that the filter responses (activation maps) produced after a convolution operation are a 4D tensor \mathbf{X} with shape $[B, W, H, C]$, where B is the mini-batch size, W, H are the spatial extents of the map, and C is the number of filters used in convolution. C is also referred to as output channels. Let $\mathbf{x} = \mathbf{X}_{b,\dots,c} \in \mathbb{R}^N$, where $N = W \times H$, be the vector of filter responses for the c^{th} filter for the b^{th} batch point. Let $\nu^2 = \sum_i x_i^2 / N$, be the mean squared norm of \mathbf{x} . Then we propose Filter Response Normalization (FRN) as following:

$$\hat{\mathbf{x}} = \frac{\mathbf{x}}{\sqrt{\nu^2 + \epsilon}}, \quad (1)$$

where ϵ is a small positive constant to prevent division by zero. A few observations are in order about the proposed normalization scheme:

1. Similar to other normalization schemes, FRN removes the scaling due to both the filter weights and pre-activations.

This is known [28] to remove noisy updates along the direction of the weights and reduce gradient covariance.

2. One of the main differences in our proposal is that we do not remove the mean prior to normalization. While mean subtraction was an important aspect of Batch Normalization, it is arbitrary and without real justification for normalization schemes that are batch independent.
3. Our normalization is done on a *per-channel* basis. This ensures that all filters (or rows of a weight matrix) have the same relative importance in the final model.
4. At a first glance, FRN appears very similar to Local Response Normalization (LRN) proposed in Krizhevsky et al. [19]. However, among other differences, LRN does normalization over adjacent channels at the same spatial location, while FRN is a global normalization over the spatial extent.

As with other schemes, we also perform an affine transform after normalization so that the network can undo the effects of the normalization:

$$\mathbf{y} = \gamma \hat{\mathbf{x}} + \beta, \quad (2)$$

where γ and β are learned parameters. The final addition to our FRN layer is the activation function.

3.1.1 Thresholded Linear Unit (TLU)

Lack of mean centering in FRN can lead to activations having an arbitrary bias away from zero. Such a bias in conjunction with ReLU can have a detrimental effect on learning and lead to poor performance and dead units. We propose to address this issue by augmenting ReLU with a learned threshold τ to yield TLU defined as:

$$\mathbf{z} = \max(\mathbf{y}, \tau) \quad (3)$$

Since $\max(\mathbf{y}, \tau) = \max(\mathbf{y} - \tau, 0) + \tau = \text{ReLU}(\mathbf{y} - \tau) + \tau$, the effect of TLU activation is the same as having a shared bias before and after ReLU. However, this does not appear to be identical to absorbing the biases in the previous and subsequent layers based on our experiments. We hypothesize that the form of TLU is more favorable for optimization. TLU significantly improves the performance of models using FRN (see Table 5), outperforming BN and other alternatives, and leads to our method, FRN layer. Figure 2 shows the schematic for our proposed FRN layer.

3.2. Gradients of FRN Layer

In this section, we derive expressions for the gradients that flow through the network in the presence of the FRN layer. Since all the transformations are performed channel-wise, we only derive the per-channel gradients. Let us assume that somewhere in the network, the activations \mathbf{x} are fed to the FRN layer and the output is \mathbf{z} (following the transformations

described in equations (1), (2), and (3)). Let $f(\mathbf{z})$ be the mapping that the network applies to \mathbf{z} , with gradients $\frac{\partial f}{\partial \mathbf{z}}$ flowing backwards. Note that the parameters γ, β and τ are vectors of size num_channels , and so the per channel updates are scalar.

$$\frac{\partial z_i}{\partial \tau} = \begin{cases} 0, & \text{if } y_i \geq \tau \\ 1, & \text{otherwise} \end{cases} \quad (4)$$

Note that the gradients $\frac{\partial z_i}{\partial y_i}$ are just the same as above, but with the cases reversed. Then the gradient update to τ is of the form

$$\frac{\partial f}{\partial \tau} = \sum_{b=1}^B \left(\frac{\partial f}{\partial \mathbf{z}_b} \right)^T \frac{\partial \mathbf{z}_b}{\partial \tau}, \quad (5)$$

where \mathbf{z}_b is the vector of per-channel activations of the b^{th} batch point. Gradients w.r.t γ and β are as follows:

$$\left(\frac{\partial f}{\partial \gamma}, \frac{\partial f}{\partial \beta} \right) = \left(\sum_{b=1}^B \frac{\partial f^T}{\partial \mathbf{y}_b} \hat{\mathbf{x}}_b, \sum_{b=1}^B \frac{\partial f}{\partial \mathbf{y}_b} \right) \quad (6)$$

Using eqn. (2), we can see that $\frac{\partial f}{\partial \hat{\mathbf{x}}} = \gamma \frac{\partial f}{\partial \mathbf{y}}$. Finally, the gradients that flow back from the FRN layer can be written as

$$\frac{\partial f}{\partial \mathbf{x}} = \frac{1}{\sqrt{\nu^2 + \epsilon}} \left(I - \frac{\hat{\mathbf{x}} \hat{\mathbf{x}}^T}{N} \right) \frac{\partial f}{\partial \hat{\mathbf{x}}} \quad (7)$$

We make a couple of observations about the gradients. Eqn. (5) suggests that part of the gradients that get suppressed in a regular ReLU activation are now used to update τ , and in some sense are *not wasted*. Since $\|\hat{\mathbf{x}}\|_2^2 = N$, eqn. (7) shows that the gradients w.r.t \mathbf{x} are orthogonal to \mathbf{x} (provided $\epsilon = 0$) because $(I - \hat{\mathbf{x}} \hat{\mathbf{x}}^T / N)$ projects out the component in the direction of $\hat{\mathbf{x}}$. This property is not unique to our normalization, but is known to help in reducing variance of gradients during SGD and benefit optimization [28].

3.3. Learning ϵ

In our discussion so far, we have assumed that the filter responses have a large spatial extent of size $N = W \times H$. However, there are situations in real networks like InceptionV3 [33] and VGG-A [30], where some layers produce 1×1 activation maps. In this setting ($N=1$), for small value of ϵ , the proposed normalization as in Equation (1) turns into a sign function (see Figure 3), and has very small gradients almost everywhere. This will invariably affect the learning adversely. In contrast, higher values of ϵ lead to variants of smoother soft sign function that are more amenable to learning. Therefore, appropriate value of ϵ becomes crucial for models that are fully connected or lead to 1×1 activation maps. Empirically, we turn ϵ into a learnable parameter (initialized at 10^{-4}) for such models. For other models, we use




Figure 3: Effect of ϵ on normalized activations for the case of $N = 1$. For very small values of ϵ , FRN turns into a step function while for higher values it behaves like a softsign function, allowing the gradients to flow. Having a learnable epsilon is crucial in models with fully connected layers or low-dimensional activation maps.

a fixed constant value of 10^{-6} . In our experiments, we show that the learnable parameterization is useful for training InceptionV3 model where the Auxiliary logits head produces 1×1 activation maps, and for the VGG-A [30] architecture that uses fully connected layers.

Since $\epsilon > 0$, we explored two alternative parameterizations to enforce this constraint: absolute value and exponential. While both trained well, the absolute value parameterization $\epsilon = 10^{-6} + |\epsilon_l|$ (ϵ_l being a learned parameter), produced consistently better empirical results. Parameterizations of this form are also preferable because the gradient magnitudes for ϵ_l are independent of the value of ϵ .

3.4. Mean Centering

Batch Normalization was proposed to counter the effects of *internal covariate shift* during training of a deep neural network. The solution was to keep the statistics of distribution of activations over the data set invariant; and as a practical matter, they choose to normalize the first and second moments of mini-batch at each step. Batch independent alternatives that include mean centering are not justified by any particular consideration, and appear merely to be a legacy of Batch Normalization.

Consider the example of Instance Normalization (IN). Using the same notation as Section 3.1, IN computes the normalized activations using the channel statistics $\mu = \sum_i x_i / N$ and $\sigma^2 = \sum_i (x_i - \mu)^2 / N$ as following:

$$\hat{\mathbf{x}} = \frac{\mathbf{x} - \mu}{\sqrt{\sigma^2 + \epsilon}} \quad (8)$$

As the size of the activation map decreases (as is common in the layers closer to the output which are subject to down-sampling, or due to the presence of fully connected layers), IN produces zero activations. Layer and Group Normalization are ways to circumvent this issue by normalizing

across (all or subset of) channels. Since individual filters are responsible for each channel’s activations, normalizing across channels introduces unnecessary interactions in the filter updates. Hence, it appears that the only principled approach is to normalize each channel of the activation map separately without resorting to mean centering. This also has the desirable effect of removing the relative scaling between filters, which is known to greatly aid in optimization.

A negative impact of not performing mean centering is that the activations can be biased arbitrarily away from zero, rendering ReLU activation less than ideal. We mitigate this issue by introducing the Thresholded Linear Unit (TLU) in Section 3.1.1. Empirically, the combination of uncentered normalization with the TLU activation outperforms BN and all other alternatives.

3.5. Implementation and Practical Considerations

FRN is easy to implement in automatic differentiation frameworks. We provide an example implementation using the python API for Tensorflow in Listing 1. Further, for FRN to achieve peak accuracy, we found that care must be exercised with the following practical considerations. Sensitivity of FRN to each of these is architecture dependent.

Learning Rate Schedule: We found that the more common step decay learning rate (LR) schedule was not optimal for FRN. Instead, continuous LR decay schedules such as cosine decay (without restarts) performed better for all methods and eliminated the need to tune step decay hyper parameters.

Warm-up: Since FRN does not perform mean centering, we empirically found that certain architectures are more sensitive to the choice of initial LR. Setting a high initial LR causes large updates, that lead to large activations in the early part of the training and result in a slowdown in the learning. This is due to the $\frac{1}{\sqrt{v^2 + \epsilon}}$ factor in the gradient of $\frac{\partial l}{\partial w}$ (see Equation (7)). This happens more often in architectures that employ several max pooling layers, like VGG-A. We address this by using an initial warp-up phase where the LR is slowly increased from 0 to the peak value. Since all our experiments use cosine LR decay schedule, we use a cosine warm-up schedule as well. Note that a warmup phase is quite common and frequently used in training [11, 13, 14].

Learnable ϵ : As discussed in Section 3.3, for models that use 1×1 activation maps, it is crucial to turn the ϵ into a learned parameter and initialize with a larger value to prevent step function like behavior and enable training.

4. Experiments

We evaluate our method extensively on two tasks: 1) Image classification on Imagenet, and 2) Object detection on COCO. While Image classification is the de-facto standard

Listing 1: Tensorflow implementation of FRN layer

```
def FRNLayer(x, tau, beta, gamma, eps=1e-6):
    # x: Input tensor of shape [BxHxWxC].
    # tau, beta, gamma: Variables of shape [1, 1, 1, C].
    # eps: A scalar constant or learnable variable.

    # Compute the mean norm of activations per channel.
    nu2 = tf.reduce_mean(tf.square(x), axis=[1, 2],
                        keepdims=True)

    # Perform FRN.
    x = x * tf.rsqrt(nu2 + tf.abs(eps))

    # Return after applying the Offset-ReLU non-linearity.
    return tf.maximum(gamma * x + beta, tau)
```

for evaluation, Object detection typically requires high resolution inputs and is particularly constrained by the large batch size requirement of BN. On Imagenet classification, we show that our method outperforms other normalization methods on three different network architectures. Further, our method does this consistently at all batch sizes we experimented with. Finally, we validate the performance of our method on Object Detection where it outperforms other normalization methods on all batch sizes as well.

4.1. ImageNet Classification

Dataset: ImageNet classification dataset [8] consists of 1000 classes. We train on the $\sim 1.28M$ training images and report results on the 50000 validation images. For all models in this section, we resize the images to 299×299 and use data augmentation from [34] at training time.

Model architectures: We provide comparisons using three different model architectures: 1) ResnetV2-50 [14]: Popular model with identity shortcuts, 2) InceptionV3 [32]: High performing model *without* identity shortcuts and fully connected layers and, 3) VGG-A [30]: Feed forward model with a mix of convolutional and fully connected layers. For all models using GN, we use a group size of 32. However, since VGG-A does not use a multiple of 32 filters in all layers, we increase the number of filters to nearest multiple.

Training: We follow the training setup used by He et al. [13]. All models are trained using synchronous SGD across 8 GPUs for 300000 steps. Gradients are computed by averaging across all GPUs. For BN, the normalization statistics are computed per GPU. This setup is common for multi-GPU training using synchronous SGD in Tensorflow and PyTorch. An initial learning rate of $0.1 \times \text{batch_size} / 256$ and cosine decay schedule is used. We follow [13, 14] for other implementation details. Results are reported using two image classification metrics: 1) ‘Precision@1’ measures the accuracy using the highest scoring class (top-1 prediction) while, 2) ‘Recall@5’ measures the accuracy using top-5 scoring classes.

Comparison with normalization methods: In Table 1 we compare our method with various normalization methods for

Table 1: **FRN layer outperforms BN and other normalization methods for large batch size on Imagenet Classification for ResnetV2-50 [14] and InceptionV3 [32].**

Method	ResnetV2 50		InceptionV3	
	P@1	R@5	P@1	R@5
Batchnorm	76.21	92.98	78.24	94.07
BatchRenorm	75.85	92.90	78.19	94.01
Groupnorm	75.67	92.70	78.14	93.98
Layernorm	72.75	91.19	76.75	93.37
Instancenorm	71.63	90.46	73.93	91.60
FRN layer [Ours]	77.21	93.57	78.95	94.49

Table 2: **Effect of mini-batch size used for normalization on ImageNet classification for ResnetV2-50 [14].**

Images per GPU →		32	16	8	4	2	1
Precision@1	Batchnorm	76.21	75.55	74.04	71.96	65.09	1.58
	Renorm	75.85	75.96	75.59	74.18	70.75	37.55
	Groupnorm	75.67	75.77	76.14	76.02	76.20	75.93
	FRN layer [Ours]	77.21	77.10	77.16	77.18	77.33	77.36
	Δ	+1.54	+1.33	+1.02	+1.16	+1.13	+1.43
Recall@5	Batchnorm	92.98	92.81	92.12	90.98	86.51	4.00
	Renorm	92.90	92.98	92.80	92.10	89.81	57.18
	Groupnorm	92.70	92.72	92.89	92.87	92.92	92.73
	FRN layer [Ours]	93.62	93.59	93.60	93.49	93.61	93.61
	Δ	+0.92	+0.87	+0.71	+0.62	+0.69	+0.88

the regular batch size of 32 images/GPU. This results in an effective batch size of $32 \times 8 = 256$ and is the most favorable configuration for BN. This is the strongest baseline for image classification and all the alternatives to BN have struggled in this setting, underperforming BN. *Even for this large batch size, FRN outperforms all the methods, including BN, with a healthy margin on both the architectures. Key takeaway is that batch dependent training is not necessary for high performance.* At this large batch size, the next best performing normalization schemes are BN and BatchRenorm, both of which are batch normalized methods, followed by other sample based normalization methods. Figure 4 compares the training and validation 'Precision@1' curves for various normalization methods using the ResnetV2-50 architecture. We observe that FRN layer achieves both higher training and validation accuracies than BN indicating that removal of stochastic batch dependence eases optimization allowing model to train better. The generalization gap, i.e. difference between training and validation accuracy, has also increased. However, improved optimization results in a net performance gain on validation. In comparison, GN also achieves lower training error than BN but performs worse on validation.

Effect of small number of images per GPU: We study the impact of mini-batch sizes used for normalization (images/GPU) on the performance of various methods in Fig-




Figure 4: **Comparison of training and validation curves of various normalization methods for Imagenet Classification using ResnetV2-50 model.**

ure 1 and Table 2. All methods are trained with 8 GPUs using six different total batch sizes of 8, 16, 32, 64, 128, 256, divided into equal number of images per GPU leading to 1, 2, 4, 8, 16, and 32 images/GPU. BN is known to degrade in performance when the batch size is small [17, 31] as evident in Figure 1. GroupNorm (GN) exhibits a more consistent performance underperforming BN only at the largest batch size. Batch renormalization outperforms GN at the largest two batch sizes but shows a degradation in performance for the smaller batch sizes. *Our method, FRN, consistently outperforms all the normalization methods at all batch sizes.*

Analyzing the effect of FRN and TLU: In Table 3 we perform a detailed ablation study of the effect of FRN and TLU. We combine them with various normalization methods – BatchNorm (BN), GroupNorm (GN), LayerNorm (LN) and InstanceNorm (IN), and train models for each combination for two high performing, but different, model architectures – ResnetV2-50 and InceptionV3. We either replace ReLU activation with TLU, or modify the normalization technique to suppress mean centering and dividing by uncentered second moments instead of variance (Equation (1) instead of Equation (8)). The corresponding normalization are named with a FRN suffix in Table 3 – for example, GN becomes GFRN, LN becomes LFRN etc. For BN, we just replaced the activation function without changing the normalizing technique, and we observe no significant difference in performance. We note, however, that IN benefits from use of FRN (IN+ReLU vs. FRN+ReLU) resulting in 3.61 P@1 gain for ResnetV2-50. Adding TLU leads to another 1.97 points gain (FRN + TLU). Similar improvements are observed for InceptionV3. In fact, similar improvement trends can be seen for GN and LN as well. This experimental result suggests that *both FRN and TLU are critical for the high performance of our method and provide complementary gains.*

Models with Fully Connected (FC) layers: FC layers are a pathological case for normalization methods, especially for

Table 3: **Ablation of our method on Imagenet Classification** for ResnetV2-50 [14] and InceptionV3 [32]. We evaluate various combinations of our method with existing normalizations. Combinations that include one of our proposals are marked as †. Our method, FRN + TLU, at the bottom is marked as [Ours].

Method	ResnetV2-50		InceptionV3	
	P@1	R@5	P@1	R@5
BN + ReLU	76.21	92.98	78.24	94.07
BN + TLU †	76.03	92.94	78.22	94.13
GN + ReLU	75.67	92.70	78.14	93.98
GN + TLU †	76.59	93.16	78.50	94.18
GFRN + ReLU †	75.93	92.65	78.16	94.03
GFRN + TLU †	76.44	92.80	78.18	94.05
LN + RELU	72.75	91.19	76.75	93.37
LN + TLU †	73.99	91.60	77.21	93.48
LFRN + RELU †	75.03	92.50	77.62	93.65
LFRN + TLU †	76.17	92.89	78.12	94.02
IN + ReLU	71.63	90.46	73.93	91.60
IN + TLU †	71.72	90.53	74.81	92.01
FRN + ReLU †	75.24	92.65	77.98	94.02
FRN + TLU [Ours]	77.21	93.57	78.95	94.49

per sample methods (GN, LN, IN, FRN), since the number of activations to be normalized over is relatively small. As a result, normalization layers are typically not applied after FC layers. In this section we evaluate the effect of applying normalization after all the layers irrespective of whether they are FC or convolutional layers. Note that FC layers are the most challenging scenario for FRN since we are normalizing over a single activation ($N = 1$). We report results for two architectures where the output of FC layers is normalized: 1) InceptionV3 in Table 1 and 2) VGG-A in Table 4. Note that while ResnetV2-50 also has a FC layer after the global pooling to produce logits, normalization is performed before pooling and is thus not relevant here. InceptionV3 has fully connected layers in an auxiliary logits branch while VGG-A has them in the main network. *FRN outperforms all other normalization methods even in this challenging scenario on both the architectures.*

Note that, while training InceptionV3 and VGG-A, it was crucial to use learning rate warm-up (refer Section 3.5) and learned ϵ (refer Section 3.3) for FRN to achieve peak performance. FRN underperformed other methods on InceptionV3 and failed to learn entirely on VGG-A without warm-up. Other methods were not significantly affected. We discovered that without the warm-up phase, the output of max pooling layers grew to very large magnitudes in first few steps. This saturates the normalized activations (see Figure 3) and prevents learning due to poor flow of gradients.

Interestingly, for VGG-A, BN performs worse than ‘No normalization’ at the default learning rate of 0.01. In Table 4 we also report results for models trained with a higher

Table 4: **Model with fully connected layer.** We provide a comparison on Imagenet Classification for the VGG-A model that uses two fully connected layers. Top half shows the results for training with an initial learning rate of 0.01 (the default rate). Bottom half shows the results for training with a higher learning rate of 0.1. The base model diverges at this rate, while the model with Batchnorm exhibits instability. FRN and Groupnorm train well, with FRN outperforming all others.

Method	Learning rate	P@1	R@5
No normalization	0.01	69.04	88.99
Batchnorm	0.01	67.82	88.11
Groupnorm	0.01	69.35	89.12
FRN	0.01	70.04	89.42
No normalization	0.1	Diverged	Diverged
Batchnorm	0.1	62.61	84.56
Groupnorm	0.1	69.94	89.57
FRN	0.1	71.66	90.69

learning rate of 0.1. A warm-up phase was useful for all the models at this learning rate. However, the ‘No normalization’ model eventually diverges, while BN shows instability in training (even with warm-up) and performs significantly worse than other methods. In contrast, both FRN and GN benefit from training at higher learning rate and yield improved performance with FRN outperforming GN.

Comparison of TLU with related variants: In Table 5 we compare TLU with three related variants for ResnetV2-50 on ImageNet. All four correspond to different combinations of having a scale κ and bias τ to compute the threshold. First observe that TLU, despite having a less general form, outperforms others. Second, all variants with a learnable threshold outperform BN, which doesn’t benefit from it. We conclude that a learnable threshold is necessary for high performance in conjunction with FRN however it doesn’t need to be input dependent. Interestingly, while two of the variants correspond to commonly known activations – ReLU and Parametric ReLU (PReLU) [12], the third more general form, termed Affine-TLU, outperforms the previous two and has not been explored to the best of our knowledge. Note that Affine-TLU is different from Maxout [10], which computes maximum across groups of channels and, unlike Affine-TLU, results in reduced number of channels.

4.2. Object Detection on COCO

Next, we evaluate our method on the task of Object Detection (OD) and demonstrate that it consistently outperforms other normalization methods at all the batch sizes we evaluated on. Since OD frameworks are typically trained with high resolution inputs, they are limited to using small mini-batch sizes. This constraint makes OD an ideal evaluation benchmark for sample based normalization methods that enable training with small batch sizes.

Table 5: **Comparison of activations** in conjunction with FRN on Imagenet Classification for ResnetV2-50. We observe that learnable threshold is key to high performance of our method in comparison to BN, which doesn't benefit from it.

Method	P@1	R@5
BN + $\max(x, 0)$ (ReLU)	76.21	92.98
BN + $\max(x, \tau)$ (TLU)	76.03	92.94
FRN + $\max(x, 0)$ (ReLU)	75.24	92.65
$\max(x, \kappa x)$ (PReLU) [12]	76.43	93.30
$\max(x, \kappa x + \tau)$ (Affine-TLU)	76.71	93.32
$\max(x, \tau)$ (TLU)	77.21	93.57

Experimental setup. We perform experiments on the COCO dataset [22] with 80 object classes. We train using the train2017 set, and evaluate on the 5k images in val2017 (minival) split. We report the standard COCO evaluation metrics of mean average precision with different IoU thresholds, namely AP, AP⁵⁰, AP⁷⁵ Lin et al. [22].

Model: We use the RetinaNet [24] object detection framework. RetinaNet is a unified single stage detector that comprises of three conceptual components: 1) A backbone network, with an off-the-shelf architecture, that acts as a convolutional feature extractor for a given high resolution input image, 2) a convolutional object classification sub-network that acts on the features extracted by the backbone network and, 3) a convolutional bounding box regression sub-network. We use a ResnetV1-101 Feature Pyramid Network backbone [23] and resize the input images to 1024×1024.

Training: To simplify experimentation and evaluation, we only compare all methods on models trained from scratch. We justify this choice based on conclusions from [15] that, by training longer, model trained from scratch can catch up with models trained by fine-tuning pre-trained models. To ensure this, we start with a baseline fine-tuned model, trained by us at the largest batch size 64, that achieves an AP of 38.3 in 25K training steps (BN*, Table 6) and is close to the corresponding result of 39.1 reported in [24]. Next, we empirically find the nearest multiple of 25K that achieves similar accuracy when training from scratch to be 125K steps (BN, Table 6). We set 125K as the base number of training steps for the largest batch size. We train our models using 8 GPUs and experiment with batch sizes in {64, 32, 16} leading to {8, 4, 2} images per GPU respectively. For smaller batch size M we set the training steps $125000 \times 64 / M$ and learning rate as $\text{base_lr} \times M / 64$. We report best performance using $\text{base_lr} \in \{0.01, 0.05, 0.1\}$. All models are trained using a momentum of 0.9 and weight decay of 4×10^{-4} .

Comparison of normalization methods: In Table 6 we observe that FRN outperforms both BN and GN at all batch

Table 6: **Object detection results on COCO.** Our method, FRN, outperforms other methods for all batch sizes. Note that while BN shows a dramatic drop in performance for smaller batch sizes, FRN exhibits a comparatively smaller degradation and consistently outperforms GN that also exhibits similarly stable performance. Note that BN* models were trained by fine-tuning a imagenet pre-trained model, while others are trained from scratch.

Method	AP			AP ⁵⁰			AP ⁷⁵		
	8	4	2	8	4	2	8	4	2
imgs/gpu									
BN*	38.3	37.1	32.9	57.2	55.4	49.1	41.5	40.4	35.9
BN	38.7	37.9	30.2	56.6	55.2	44.5	42.1	41.4	32.5
GN	39.3	39.0	38.7	57.8	57.5	56.9	42.6	42.3	41.8
FRN	39.6	39.5	39.1	58.5	58.4	57.5	43.1	43.3	42.3

sizes, further validating our results in the previous section. In agreement with the observations from Table 2 both FRN and GN achieve higher accuracy than BN at the evaluated batch sizes. FRN outperforms BN by a significant difference of **0.9** AP points at the largest batch size, and this gap widens to **8.9** AP points at the smallest batch size. Further, FRN consistently achieves higher accuracy than GN.

Effect of batch size: BN exhibits a dramatic degradation in performance, dropping by **8.5** AP points for the model trained from scratch, as the number of images per GPU is reduced to 2. In comparison, both FRN and GN show a relatively more stable accuracy and degrade by less than **0.6** AP points. Interestingly, the finetuned BN* model for the smallest batch size performs 2.7 AP points better than the corresponding BN model trained from scratch, indicating that longer training at this batch size is detrimental to the performance of batchnorm. In contrast, FRN maintains a consistent lead for all the metrics across all batch sizes.

5. Conclusion

In this paper we proposed the FRN layer, a novel combination of Filter Response Normalization (FRN) and a Thresholded activation (TLU) function that eliminates the need for batch dependent training. It outperforms BN in a variety of settings and exhibits a consistently high performance in large as well as small batch training. Further, FRN also outperforms Group Normalization, a leading sample based normalization alternative to BN, in all the explored settings. We also demonstrated the success of FRN in the pathological case of fully connected layers which are typically not normalized. However, since different normalization methods have been successful in different problem domains, e.g. Layer Normalization has been successful in NLP, we leave exploration of these areas with FRN as future work.

Acknowledgement. We would like to thank Vivek Rathod for help with object detection experiments.

References

- [1] Devansh Arpit, Yingbo Zhou, Bhargava U Kota, and Venu Govindaraju. Normalization propagation: A parametric technique for removing internal covariate shift in deep networks. *arXiv preprint arXiv:1603.01431*, 2016. [3](#)
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. [1](#), [2](#)
- [3] Johannes Ballé, Valero Laparra, and Eero Simoncelli. Density modeling of images using a generalized normalization transformation. In *ICLR*, 2016. [3](#)
- [4] Johannes Ballé, Valero Laparra, and Eero Simoncelli. End-to-end optimized image compression. In *ICLR*, 2017. [3](#)
- [5] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. *ICLR*, 2018. [3](#)
- [6] AB Bonds. Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. *Visual neuroscience*, 2(1):41–55, 1989. [3](#)
- [7] Matteo Carandini and David J Heeger. Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13(1):51, 2012. [3](#)
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition*, pages 248–255. Ieee, 2009. [5](#)
- [9] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256, 2010. [2](#)
- [10] Ian J Goodfellow, David Warde-Farley, Mehdi Mirza, Aaron Courville, and Yoshua Bengio. Maxout networks. *arXiv preprint arXiv:1302.4389*, 2013. [7](#)
- [11] Priya Goyal, Piotr Dollár, Ross B. Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch SGD: training imagenet in 1 hour. *CoRR*, abs/1706.02677, 2017. URL <http://arxiv.org/abs/1706.02677>. [5](#)
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015. [7](#), [8](#)
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Computer vision and pattern recognition*, pages 770–778, 2016. [5](#)
- [14] 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. [5](#), [6](#), [7](#)
- [15] Kaiming He, Ross Girshick, and Piotr Dollár. Rethinking imagenet pre-training. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4918–4927, 2019. [8](#)
- [16] David J Heeger. Normalization of cell responses in cat striate cortex. *Visual neuroscience*, 9(2):181–197, 1992. [3](#)
- [17] Sergey Ioffe. Batch renormalization: Towards reducing mini-batch dependence in batch-normalized models. In *Advances in Neural Information Processing Systems*, pages 1942–1950, 2017. [1](#), [2](#), [6](#)
- [18] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *CoRR*, abs/1502.03167, 2015. [1](#), [2](#)
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’12, pages 1097–1105, USA, 2012. Curran Associates Inc. URL <http://dl.acm.org/citation.cfm?id=2999134.2999257>. [3](#)
- [20] Yann LeCun, Léon Bottou, Genevieve B Orr, and Klaus-Robert Müller. Efficient backprop. In *Neural networks: Tricks of the trade*, pages 9–50. Springer, 1998. [2](#)
- [21] Yann A LeCun, Léon Bottou, Genevieve B Orr, and Klaus-Robert Müller. Efficient backprop. In *Neural networks: Tricks of the trade*, pages 9–48. Springer, 2012. [2](#)
- [22] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. [8](#)
- [23] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. [8](#)
- [24] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. [8](#)
- [25] Chao Peng, Tete Xiao, Zeming Li, Yuning Jiang, Xiangyu Zhang, Kai Jia, Gang Yu, and Jian Sun. Megdet: A large mini-batch object detector. In *Computer Vision and Pattern Recognition*, pages 6181–6189, 2018. [2](#)
- [26] Mengye Ren, Renjie Liao, Raquel Urtasun, Fabian H Sinz, and Richard S Zemel. Normalizing the normalizers: Comparing and extending network normalization schemes. *arXiv preprint arXiv:1611.04520*, 2016. [3](#)
- [27] John H Reynolds and David J Heeger. The normalization model of attention. *Neuron*, 61(2):168–185, 2009. [3](#)
- [28] Tim Salimans and Diederik P Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In *Advances in Neural Information Processing Systems*, pages 901–909, 2016. [3](#), [4](#)
- [29] Eero P Simoncelli and David J Heeger. A model of neuronal responses in visual area mt. *Vision research*, 38(5):743–761, 1998. [3](#)
- [30] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. [4](#), [5](#)
- [31] Saurabh Singh and Abhinav Shrivastava. Evalnorm: Estimating batch normalization statistics for evaluation. In *ICCV*, 2019. [1](#), [2](#), [6](#)
- [32] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception ar-

- chitecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016. 5, 6, 7
- [33] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*,, 2016. URL <http://arxiv.org/abs/1512.00567>. 4
- [34] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *AAAI*, volume 4, page 12, 2017. 5
- [35] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016. 2
- [36] Yuxin Wu and Kaiming He. Group normalization. In *Computer Vision and Pattern Recognition*, 2018. 1, 2