

High Frequency Residual Learning for Multi-Scale Image Classification

Bowen Cheng¹
bcheng9@illinois.edu

Rong Xiao²
rongxiao@gmail.com

Jianfeng Wang³
jianfw@microsoft.com

Thomas Huang¹
t-huang1@illinois.edu

Lei Zhang³
leizhang@microsoft.com

¹ University of Illinois at
Urbana-Champaign
Urbana, IL, USA

² Ping An Property&Casualty Insurance
Company of China
China

³ Microsoft
Redmond, WA, USA

Abstract

We present a novel high frequency residual learning framework, which leads to a highly efficient multi-scale network (MSNet) architecture for mobile and embedded vision problems. The architecture utilizes two networks: a low resolution network to efficiently approximate low frequency components and a high resolution network to learn high frequency residuals by reusing the upsampled low resolution features. With a classifier calibration module, MSNet can dynamically allocate computation resources during inference to achieve a better speed and accuracy trade-off. We evaluate our methods on the challenging ImageNet-1k dataset and observe consistent improvements over different base networks. On ResNet-18 and MobileNet with $\alpha = 1.0$, MSNet gains 1.5% accuracy over both architectures without increasing computations. On the more efficient MobileNet with $\alpha = 0.25$, our method gains 3.8% accuracy with the same amount of computations.

1 Introduction

Convolutional neural networks (CNNs) have become the dominant machine learning approach in solving computer vision problems such as, image classification [9, 12, 13, 24, 36, 38], object detection [10, 11, 16, 17, 32, 43], semantic segmentation [3, 4, 5, 26, 50], etc. In the past several years, great progresses have been achieved in the study of developing large and computationally intensive networks [18, 36, 39] to achieve higher accuracy with the cost of sacrificing speed and efficiency.

Meanwhile, there have been increasing needs of deploying convolutional neural networks into mobile devices with limited computational resources. To address the computational efficiency problem, works [20, 34, 50] have been proposed to discover more efficient network structure. Common strategies such as replacing convolution operation with more efficient

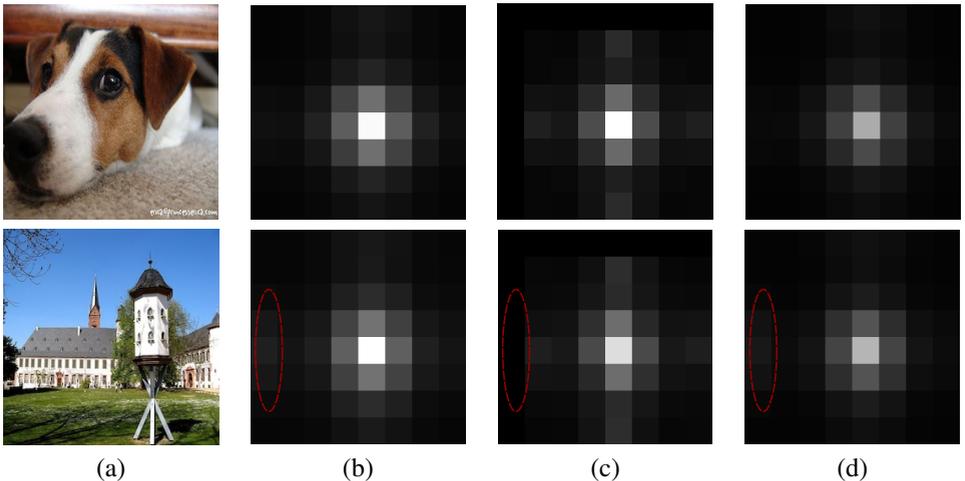


Figure 1: Frequency domain analysis. We visualize the magnitude of 2D discrete Fourier transform of final feature map of different networks. We plot the average magnitude of all feature maps from the last convolutional layer. (a): input image to the network. (b): mean magnitude for ResNet-18, it contains mostly the low frequency component. (c): mean magnitude for low resolution network of MS-ResNet-18, compared with ResNet-18, its low frequency component is more concentrated to the DC component. (d): mean magnitude for high resolution network of MS-ResNet-18, compared with ResNet-18 and its low resolution counterpart, it has far less low frequency component. The red ellipse partially covers the high frequency region. From ResNet-18, it shows strong high frequency magnitude. From lower resolution network of MS-ResNet-18, the covered frequency vanishes; while the higher resolution network recovers the high frequency.

group or depth-wise separable convolutions, using less channels and using lower resolution input have been used.

Among these strategies, reducing the input resolution by half is probably the simplest method with the advantage that it can be directly applied during inference. Without the need of re-training, it can reduce the computation cost quadratically. However, while it decreases the computational cost, it usually sacrifices the overall accuracy. One question we want to ask and address is: can we do inference on low resolution images without losing accuracy? The answer is yes for some images but not for all the images. Recent works [0, 8, 25, 42] propose to directly inference from low resolution input by integrating a super-resolution network with a classification network, but the gain is still marginal. We observe that there are some easy images that can be predicted correctly using both their low resolution and high resolution versions (Figure 2 (a)). A straightforward solution is to predict a low resolution image first. If the prediction is incorrect, we make another prediction using the high resolution image. This straightforward solution has a problem that we do not know when classifier fails on low resolution. However, based on the property of softmax function, we can make an assumption that when the prediction of a classifier has a high softmax score, it is more likely to be a correct prediction (Figure 2 (b)). Based on this assumption, we can develop a classifier calibration module: we first run prediction on low resolution images, if the softmax score of the prediction is higher than a threshold, we keep this prediction; if the softmax score of the prediction is lower than the threshold, we use the prediction on the high resolution

image. Classifier calibration can be used dynamically during inference to save the average computation.

However, classifier calibration only solves the problem partially. While for an easy image we can use the low resolution predictor to save computational cost quadratically, for a hard image we need to use two predictors corresponding to low and high resolutions as we need the prediction score from the low resolution predictor to tell whether it is an easy or hard image. Checking this process closer, we can easily find that this naive solution does not reuse features of low resolution images anymore once it finishes the prediction on low resolution images, which leads to a waste of computation.

Therefore, the second question we want to ask and address is more intriguing: can low resolution features help high resolution predictions? This question is more interesting because reusing low resolution features will save more computational overheads.

Inspired by the idea of wavelet transform and residual learning, we propose a novel network structure named High Frequency Residual Multi-Scale Network (MSNet) aiming at learning low frequency components and high frequency residuals separately. Note that a feature map normally contains both low frequency components and high frequency components (Figure 1 (b)). If we upsample a low resolution feature map, it contains mostly low frequency components (Figure 1 (c)). We find that by reusing the upsampled low resolution features, the network has the ability to learn high frequency residuals with less low frequency components (Figure 1 (d)). Formally, denoting the desired high resolution feature map by \mathbf{y}_H , we hope a lower scale network (a network takes as input low resolution images) learns mainly low frequency components $\mathbf{y}_L = \mathcal{L}(\mathbf{x}_L)$ and a higher scale network (a network takes as input high resolution images) learns mainly high frequency residuals $\mathcal{H}(\mathbf{x}_H)$. Then the combination of the upsampled low frequency components and high frequency residuals is the desired complete feature map $\mathbf{y}_H = \mathcal{H}(\mathbf{x}_H) + u(\mathbf{y}_L)$. Combining *classifier calibration* and *high frequency residual learning*, we achieve a much better computation and accuracy trade-off on multiple efficient networks.

Our contributions are threefold. 1) We propose the MSNet based on high frequency residual learning to efficiently reuse multi-scale features. 2) We propose a classifier calibration module that can dynamically allocate computations during inference. 3) Our proposed network achieves consistent gains over various architectures without increasing the amount of the computation.

2 Related Works

Efficient Networks. Much attention has been placed on efficient network design recently. Most common methods for designing efficient network or efficient inference are: to use more efficient model component, *e.g.* group convolutions or even depth-wise convolutions [20, 32, 50]; to perform model compression to reduce parameters and computations [24, 47, 48]; to use quantization to reduce float32 operations to float16, integer or even binary operations [30]; to use knowledge distillation to transfer knowledge of a large network or an ensemble of networks to a small network [19, 29, 33]. In this work, we focus on the efficient network learning, especially, we target at reducing the computation by reducing the input resolution but preserve the accuracy at the same time. The other model compression/acceleration methods are orthogonal to our method and can further boost the performance.

Residual Learning. The idea of ResNet [18] is to add the bypass connection between components in a feed-forward network. It makes deeper networks easier to optimize and has been applied in a wide variety of works [13, 30, 39, 45]. [6] shows the equivalence of residual net-

works [18] and densely connected networks [21] and the authors propose a general form of residual learning with enhanced performance. Compared with ResNet, we explicitly focus on learning the high frequency residual information in a multi-scale network architecture, which allows us to easily allocate the computational cost to trade-off the accuracy. Another perspective is that ResNet learns residual in the image domain while our method targets to learn residual in the frequency domain.

Cascade Classifiers. Cascade classifiers have been widely use for efficient inference. The high level idea of cascade classifiers to classify an image with a sequence of classifiers where the earlier classifiers reject easier images which is called an “early exit”. The Viola Jones Algorithm [41] for face detection uses a hard cascade by Adaboost [15], where multiple weak classifiers take different features in a cascade manner. If any of the classifier in the sequence rejects a region, then the region will be classified as a non-facial region. A soft cascade classifier [2] builds each weak classifier based on the output of all previous classifiers. In the context of deep learning, multiple works have studied the problem of “early exit” in a deep model [11, 22, 40]. The idea is to inference using early stage features based on the computation budget, however, the problem is early stage features usually do not have high level semantics. In our work, we use the idea of “early exit”, but instead of building weak classifiers at early stage features, we build a weak classifier on a low resolution input. In this way, our method achieves a good trade-off between speed and high-level features.

Multi-Scale Network Design. Network designs, either by human [18, 21, 36] or by architecture search [31, 52], mainly focus on finding more powerful building blocks, *e.g.* residual blocks in [18], densely connected blocks in [21] and machine searched blocks in [31, 52]. However, all these works use a single resolution input (*e.g.* 224×224 or 299×299), we find less works study the effect of using features with different resolutions in classification. Recently, there are works [22, 27, 28, 37, 44, 46, 49] focusing on multi-scale design. Although MSDNet [22] claims to be “multi-scale”, we find it is more similar to [35] which is a wrapping of different networks into a single network. As far as we know, our work is one of the first to show that reusing low resolution feature helps high resolution prediction. We will give a more detailed comparison in the next section.

3 High Frequency Residual Learning

3.1 Analysis of Multi-Resolution Inference

Nowadays, most CNN model architecture are fully convolutional with a global average pooling as the last feature layer, making it straightforward to do inference at arbitrary resolution. Given a well-trained model, we can evaluate its inference accuracy over different scales. Figure 2 shows the experiment results. The dataset is ImageNet 2012 [42], and the model¹ is trained and tested on the standard train split and val split, respectively. If the test image size is 128×128 , 56% of the images could be correctly predicted, within which a majority ($52\%/56\% = 93\%$) of the images can also be correctly classified if the test size is 256×256 . Among the incorrectly predicted images by 128×128 , $19\%/44\% = 43\%$ of the images can be corrected by the larger input size of 256×256 .

In single scale inference, 128×128 input leads to lower computing cost, but only achieves 56% accuracy, while 256×256 can achieve a higher accuracy but with a penalty of higher

¹Only for the verifying purpose, we use ResNet-18 [18] trained with 224×224 crop. Otherwise, we default the crop size to 256×256 if not specified explicitly.

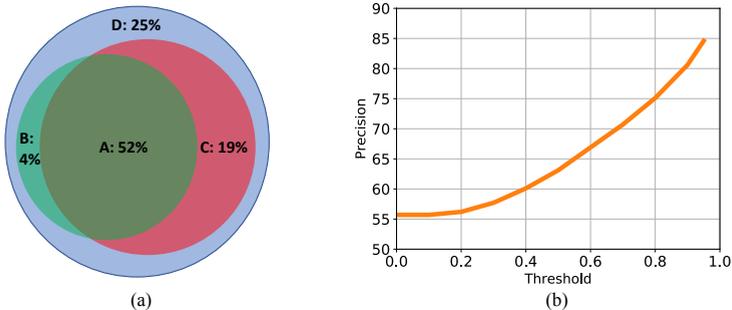


Figure 2: Multi-resolution inference results. The model is ResNet-18. (a): accuracy distribution from the input of 128×128 and 256×256 . Region A: correctly predicted by both 128×128 and 256×256 ; Region B: correctly predicted by 128×128 but incorrectly predicted by 256×256 ; Region C: incorrectly predicted by 128×128 but correctly predicted by 256×256 ; Region D: incorrectly predicted by both 128×128 and 256×256 . (b): precision vs thresholds. The threshold is used to select images to be predicted using low resolution input (128×128), precision is the accuracy of the selected image. The higher the threshold, the less images are selected and the accuracy becomes higher.

cost. If we apply the traditional multi-scale inference, where each image is tested twice, the accuracy could be even higher, but leads to a potential waste of computation. For example, 56% of the images can be handled well by using only smaller scale and testing with both scales gives little benefit (4%).

This motivates us to apply the network first onto the small scale to save the cost in the target of the 56% correctly predicted images. If the smaller scale fails, we resort to the higher scale prediction. Ideally, the upper bound of the accuracy is $56\% + 19\% = 75\%$.

The challenging problem is how to properly decide when the lower input size fails. In this paper, we adopt a simple strategy to apply a threshold on the maximum value of the probability output of the network. The intuition is that if the network is confident on the prediction (the probability is larger than a pre-defined threshold), the accuracy should also be high, which is verified in Figure 2 (b), illustrating the monotonic increasing relation of accuracy over the thresholds. The accuracy is calculated over the samples whose highest prediction score is larger than the threshold. As the threshold increases, the number of predictions decreases but their accuracy increases. When the threshold is 0, all the predictions are handled by the lower input scale. If the threshold is 1, effectively all the images are evaluated by the larger input scale.

The next question is how to learn the network such that the lower scale and the higher scale can be cooperated efficiently. A naive solution is to learn two networks independently. One is for the smaller scale, while the other is for the larger scale. However, this is inferior because the features learned for the smaller scale network are discarded in the feature learning for the larger scale network. Instead of learning two networks independently, we propose a *high frequency residual* building block to jointly learn a multi-scale network.

3.2 High Frequency Residual Learning

Figure 3 (b) depicts the proposed building block. The input is \mathbf{x}_L and \mathbf{x}_H in the real domain for lower scale and higher scale, respectively. Passing the two signals each through several

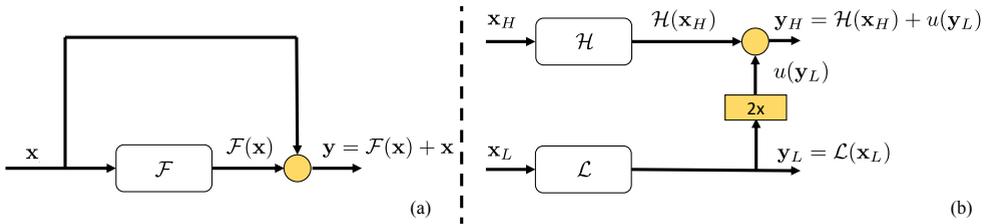


Figure 3: Residual learning block vs. high frequency residual learning block. (a) a building block for residual learning, \mathbf{x} is the input feature, $\mathcal{F}(\mathbf{x})$ is the learned residual and $\mathbf{y} = \mathcal{F}(\mathbf{x}) + \mathbf{x}$ is the output feature. (b) a building block for high frequency residual learning in Equation 1.

linear layers interleaved with nonlinear activations, we can have more abstract feature representations of $\mathcal{L}(\mathbf{x}_L)$, and $\mathcal{H}(\mathbf{x}_H)$. Instead of learning them separately, we enrich the larger scale pass by the upsampled lower signal $\mathbf{y}_L = \mathcal{L}(\mathbf{x}_L)$. That is, the output of the higher scale is $\mathbf{y}_H = \mathcal{H}(\mathbf{x}_H) + u(\mathbf{y}_L)$, where $u(\cdot)$ is an upsampling function. Formally, the building block can be written as

$$\mathbf{y}_L = \mathcal{L}(\mathbf{x}_L) \quad (1)$$

$$\mathbf{y}_H = \mathcal{H}(\mathbf{x}_H) + u(\mathbf{y}_L). \quad (2)$$

The upsampling operation helps match the spatial resolution between two different scales to make the aggregation feasible. Meanwhile, the function is a low-pass filter, which reduces the high-frequency information to pass through. For example, if we implement the upsampling by a nearest neighbor interpolation, the impulse response is a rectangle function and the corresponding frequency function is a sinc function, which allows more lower frequency to pass and blocks more higher frequency.

Since the flowing information from the bottom to the top is mostly the low frequency information, the higher scale is expected to focus on the high frequency information learning. This intuition can also be verified in Figure 1 (d), which contains less lower frequency information than 1 (b). Thus, we call it *high frequency residual learning*.

Figure 3 (a) shows the residual learning [18] for comparison. Since we explicitly split the signal into different frequency bands, it is easier to cooperatively apply the multi-scale inference efficiently. Besides, the residual learning block can be part of the high frequency learning block by inserting it into \mathcal{H} and \mathcal{L} .

3.3 High Frequency Residual Multi-Scale Network

By stacking multiple high frequency residual learning blocks, we construct a multi-scale network (MSNet) as shown in Figure 4 (a). It consists of a lower scale network, a higher scale network and a calibration module. The raw image serves as the input to the higher scale network, while the downsampled version to the smaller scale network. We use the superscript to distinguish different blocks. Within the i -th block, let \mathbf{x}_L^i and \mathbf{x}_H^i be the input to the lower scale and higher scale networks, respectively. Within each feature extraction module \mathcal{L}^i and \mathcal{H}^i , we first apply a convolutional layer with stride as 2 to reduce the feature map size. Instead of using \mathbf{y} in Eq. 1, we use \mathbf{x}_L^{i+1} and \mathbf{x}_H^{i+1} to denote outputs, which are also the input of the next block. After the last block, a global average pooling layer is added before we apply a linear layer and the softmax layer to output the classification result.

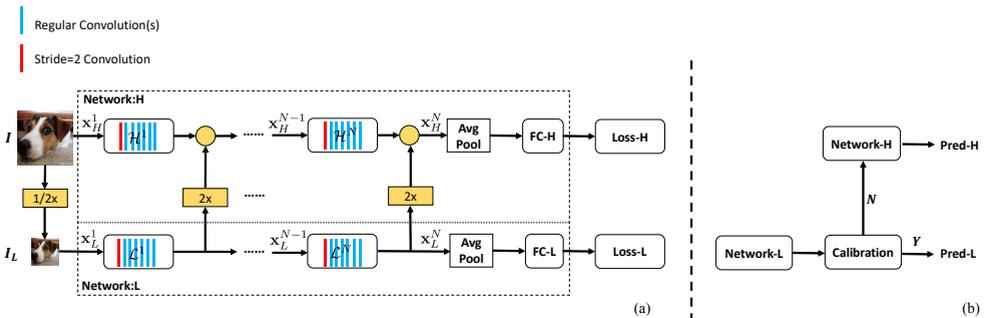


Figure 4: Multi-Scale Network (MSNet) network structure. (a) MSNet structure for training: MSNet is composed of a higher scale network (Network-H in upper dashed box) and a lower scale network (Network-L in lower dashed box). Both networks can make prediction and are trained jointly. (b) MSNet with calibration in inference: if the calibration module produces “Y” (softmax score higher than the threshold), MSNet only use Network-L prediction (pred-L); if the calibration module produces “N” (softmax score lower than the threshold), MSNet use Network-H prediction (pred-H). Note that when using Network-H, features in Network-L are reused and passed into Network-H together with the input image.

The inference stage of MSNet is shown in Figure 4 (b), the test image is first downsampled and then fed into the lower scale network. The classification result is sent to the calibration module. If the maximum probability is larger than a pre-defined threshold, the calibration module outputs this result and terminates the inference process. Otherwise, it enables the upsampling of low resolution features from the lower scale network and the forward pass of the higher scale network. In this case, the calibration module takes the output from the higher scale network as the final prediction.

During training, we remove the calibration module and add the cross entropy loss to each of the classifier. Since all the components are differentiable, we can apply the off-the-shelf gradient descent algorithm to update the parameters.

4 Experiments

4.1 Settings

We evaluate our method on the widely-used ImageNet 2012 classification dataset [14]. All the models are trained on the 1.28 million training images, and we report the top-1 accuracy of a single center crop on the 50k validation set. If the crop size is 224×224 , we first resize the shorter side of the image to 256 before doing the center crop. If the crop size is 256×256 , we resize the shorter side to 300. We use standard data augmentation in [18, 20] and train the model for 120 epochs with initial learning rate 0.1, reduced by 0.1 at 30, 60, 90 epochs.

Given a base model, e.g. ResNet, to construct its MSNet version, we extract each stage of the base model and use it as the feature abstraction part of \mathcal{H} and \mathcal{L} . We conduct experiments on ResNet-18 and MobileNet, and the corresponding MSNet version is called MS-ResNet-18 and MS-MobileNet, respectively. For the calibration module, we use different threshold to trade-off the computation and the accuracy. The computation cost is measured by the FLOPS as in [18, 20]. In MSNet, different images could have different FLOPS cost, and we report the average cost. Without explicitly explaining, the higher scale is 256×256 and the

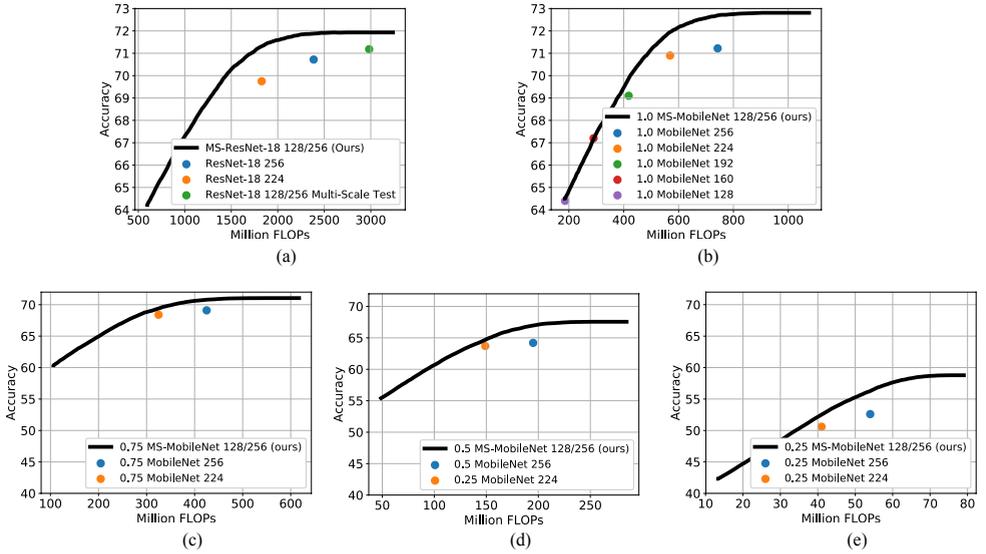


Figure 5: Our MSNet calibration results with different networks. (a) MS-ResNet-18 vs. ResNet-18. (b)-(e) MS-MobileNet vs. MobileNet. The black curves (our results) are obtained by varying threshold values.

lower scale is 128×128 .

4.2 Results

4.2.1 MS-ResNet-18 vs ResNet-18

Figure 5 shows the experiment results. Detailed numbers are listed in the supplementary material. First, we reproduce the ResNet-18 network and get an accuracy of 69.8% over the 224×224 input. By setting the threshold as 0.48, our MS-ResNet-18 can achieve similar accuracy of 69.9%, but consumes only 77% (1413/1827) FLOPs. If the threshold is 0.66, the computation cost is similar, but the accuracy of our MS-ResNet-18 is boosted to 71.3% with an absolute gain of 1.5%.

If the input size of ResNet-18 is 256×256 , the accuracy is 70.7%. Still, with similar accuracy, the FLOPs cost can be reduced to 68% (1625/2386) with our MS-ResNet-18.

Next, we perform the traditional multi-scale testing by averaging prediction results of different resolution (128 and 256) inputs for ResNet-18, which achieves 71.2% accuracy but consumes 25% more computations. Using our proposed MS-ResNet-18, we achieve 71.9% accuracy and consumes only 79% (2358/2983) computations.

4.2.2 MS-MobileNet vs MobileNet

Following [20], we shrink the number of channels by a factor of α to achieves different trade-offs on the MobileNet. The corresponding MSNet version is also shrunk accordingly. We pre-fix the network name with α in the experiment results. Figure 5 shows the experiment results. Detailed numbers are listed in the supplementary material due to limited space. For MobileNet with $\alpha = 1.0$ where α is the width multiplier to thin a network uniformly at each layer, we change the input scale to obtain different accuracies. As the input resolution

Table 1: Experiment results between our MS-DenseNet-121 and DenseNet-121.

Method	Billion FLOPs	Top-1 Accuracy
DenseNet-121 (224)	2.8	75.0%
MSDNet [24]	1.7	75.0%
MS-DenseNet-121 (112/224), thresh=0.60	1.7	75.0%
MS-DenseNet-121 (112/224), thresh=0.95	2.8	76.1%

increases, our MS-MobileNet shows larger advantages under the same computational cost over the MobileNet. This clearly demonstrates that the lower scale in our MS-MobileNet could identify correctly a certain amount of images without resorting to the higher scale network.

When α becomes smaller, we observe that our MS-MobileNet achieves even higher accuracy relatively under the same amount of computations, e.g. comparing the MobileNet with 256×256 input, the gain is 1.7% for $\alpha = 0.75$, 2.8% for $\alpha = 0.5$ and 3.8% for $\alpha = 0.25$. This clearly demonstrates our approach shows superior advantages for small models and promising applications in mobile and embedded devices.

4.2.3 MS-DenseNet-121 vs DenseNet-121

When applying our method to DenseNet, we use crop 224×224 as input and use cropping or zero-padding to align upsampled low resolution feature maps to the corresponding high resolution feature maps if their spatial resolution does not match exactly.

Table 1 shows our experimental results on DenseNet-121. The original DenseNet-121 has an accuracy of 75% with 2.8 Billion floating point operations. By setting the threshold to 0.6, our MS-DenseNet-121 achieves similar accuracy (75%) but with less computation (1.7 Billion FLOPs). That is, when combined with DenseNet, our method is on par with state-of-the-art method [24]. With the same computation (2.8 BFLOPs), our method achieves a much higher accuracy (76.1%) than the original DenseNet with a margin of 1.1%.

4.3 Speed-Accuracy Trade-off

Since different thresholds lead to different computation costs and accuracies, we can achieve a speed-accuracy trade-off by varying the threshold values. We illustrate the relationship by enumerating multiple thresholds in Figure 5 for MS-ResNet-18 and MS-MobileNet (the black curve). The baseline approaches (colored dots) are also shown in the figure. We can easily observe the advantages of the MSNet over the baselines.

5 Conclusions

In this paper, we have presented a novel high frequency residual learning framework that decouples the learning of low frequency feature and high frequency feature. We demonstrate that computations can be saved by using a low resolution network to approximate the process of learning low frequency features. We have also proposed a classifier calibration module which can dynamically allocate computation resources during inference and lead to a better speed and accuracy trade-off. In our future work, we will design more flexible high frequency residual networks that can take as input arbitrary size of images. We will also demonstrate the effectiveness of high frequency residual learning in other recognition task, e.g. object detection and semantic segmentation, in the future.

References

- [1] Tolga Bolukbasi, Joseph Wang, Ofer Dekel, and Venkatesh Saligrama. Adaptive neural networks for efficient inference. *ICML*, 2017.
- [2] Lubomir Bourdev and Jonathan Brandt. Robust object detection via soft cascade. In *IEEE CVPR*, volume 2, pages 236–243. IEEE, 2005.
- [3] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- [4] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE TPAMI*, 40(4):834–848, 2018.
- [5] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018.
- [6] Yunpeng Chen, Jianan Li, Huaxin Xiao, Xiaojie Jin, Shuicheng Yan, and Jiashi Feng. Dual path networks. In *NIPS*, pages 4470–4478, 2017.
- [7] Bowen Cheng, Zhangyang Wang, Zhaobin Zhang, Zhu Li, Ding Liu, Jianchao Yang, Shuai Huang, and Thomas S Huang. Robust emotion recognition from low quality and low bit rate video: A deep learning approach. In *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 65–70. IEEE, 2017.
- [8] Bowen Cheng, Ding Liu, Zhangyang Wang, Haichao Zhang, and Thomas S Huang. Visual recognition in very low-quality settings: Delving into the power of pre-training. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [9] Bowen Cheng, Yunchao Wei, Honghui Shi, Shiyu Chang, Jinjun Xiong, and Thomas S Huang. Revisiting pre-training: An efficient training method for image classification. *arXiv preprint arXiv:1811.09347*, 2018.
- [10] Bowen Cheng, Yunchao Wei, Honghui Shi, Rogerio Feris, Jinjun Xiong, and Thomas Huang. Decoupled classification refinement: Hard false positive suppression for object detection. *arXiv preprint arXiv:1810.04002*, 2018.
- [11] Bowen Cheng, Yunchao Wei, Honghui Shi, Rogerio Feris, Jinjun Xiong, and Thomas Huang. Revisiting rcnn: On awakening the classification power of faster rcnn. In *ECCV*, 2018.
- [12] Bowen Cheng, Rong Xiao, Yandong Guo, Yuxiao Hu, Jianfeng Wang, and Lei Zhang. Revisit multinomial logistic regression in deep learning: Data dependent model initialization for image recognition. *arXiv preprint arXiv:1809.06131*, 2018.
- [13] François Chollet. Xception: Deep learning with depthwise separable convolutions. *arXiv preprint*, 2016.
- [14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE CVPR*, pages 248–255. IEEE, 2009.

- [15] Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55 (1):119–139, 1997.
- [16] Ross Girshick. Fast r-cnn. In *IEEE ICCV*, pages 1440–1448. IEEE, 2015.
- [17] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *IEEE CVPR*, pages 580–587, 2014.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE CVPR*, pages 770–778, 2016.
- [19] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [20] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- [21] Gao Huang, Zhuang Liu, Kilian Q Weinberger, and Laurens van der Maaten. Densely connected convolutional networks. In *IEEE CVPR*, volume 1, page 3, 2017.
- [22] Gao Huang, Danlu Chen, Tianhong Li, Felix Wu, Laurens van der Maaten, and Kilian Q Weinberger. Multi-scale dense networks for resource efficient image classification. *ICLR*, 2018.
- [23] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and < 0.5 mb model size. *arXiv preprint arXiv:1602.07360*, 2016.
- [24] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012.
- [25] Ding Liu, Bowen Cheng, Zhangyang Wang, Haichao Zhang, and Thomas S Huang. Enhance visual recognition under adverse conditions via deep networks. *arXiv preprint arXiv:1712.07732*, 2017.
- [26] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *IEEE CVPR*, pages 3431–3440, 2015.
- [27] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European Conference on Computer Vision*, pages 483–499. Springer, 2016.
- [28] Marco Pedersoli, Andrea Vedaldi, Jordi Gonzalez, and Xavier Roca. A coarse-to-fine approach for fast deformable object detection. *Pattern Recognition*, 48(5):1844–1853, 2015.
- [29] Antonio Polino, Razvan Pascanu, and Dan Alistarh. Model compression via distillation and quantization. *arXiv preprint arXiv:1802.05668*, 2018.

- [30] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *ECCV*, pages 525–542. Springer, 2016.
- [31] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. 2018.
- [32] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, pages 91–99, 2015.
- [33] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014.
- [34] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Inverted residuals and linear bottlenecks: Mobile networks for classification, detection and segmentation. *arXiv preprint arXiv:1801.04381*, 2018.
- [35] Shreyas Saxena and Jakob Verbeek. Convolutional neural fabrics. In *Advances in Neural Information Processing Systems*, pages 4053–4061, 2016.
- [36] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [37] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. *arXiv preprint arXiv:1902.09212*, 2019.
- [38] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, Andrew Rabinovich, Jen-Hao Rick Chang, et al. Going deeper with convolutions. In *IEEE CVPR*.
- [39] 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.
- [40] Surat Teerapittayanon, Bradley McDanel, and HT Kung. Branchynet: Fast inference via early exiting from deep neural networks. In *Pattern Recognition (ICPR), 2016 23rd International Conference on*, pages 2464–2469. IEEE, 2016.
- [41] Paul Viola and Michael J Jones. Robust real-time face detection. *IJCV*, 57(2):137–154, 2004.
- [42] Zhangyang Wang, Shiyu Chang, Yingzhen Yang, Ding Liu, and Thomas S Huang. Studying very low resolution recognition using deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4792–4800, 2016.
- [43] Yunchao Wei, Zhiqiang Shen, Bowen Cheng, Honghui Shi, Jinjun Xiong, Jiashi Feng, and Thomas Huang. Ts2c: Tight box mining with surrounding segmentation context for weakly supervised object detection. In *ECCV*, 2018.

- [44] Guotian Xie, Jingdong Wang, Ting Zhang, Jianhuang Lai, Richang Hong, and Guojun Qi. Interleaved structured sparse convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8847–8856, 2018.
- [45] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *IEEE CVPR*, pages 5987–5995. IEEE, 2017.
- [46] Wei Yang, Shuang Li, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang. Learning feature pyramids for human pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1281–1290, 2017.
- [47] Jiahui Yu and Thomas Huang. Universally slimmable networks and improved training techniques. *arXiv preprint arXiv:1903.05134*, 2019.
- [48] Jiahui Yu, Linjie Yang, Ning Xu, Jianchao Yang, and Thomas Huang. Slimmable neural networks. *arXiv preprint arXiv:1812.08928*, 2018.
- [49] Wei Zhang, Gregory Zelinsky, and Dimitris Samaras. Real-time accurate object detection using multiple resolutions. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8. IEEE, 2007.
- [50] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. *arXiv preprint arXiv:1707.01083*, 2017.
- [51] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017.
- [52] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *IEEE CVPR*, 2018.