

Resolution Switchable Networks for Runtime Efficient Image Recognition

Yikai Wang^{1*}, Fuchun Sun¹, Duo Li², and Anbang Yao²

¹ Beijing National Research Center for Information Science and Technology (BNRist),
State Key Lab on Intelligent Technology and Systems,
Department of Computer Science and Technology, Tsinghua University

² Cognitive Computing Laboratory, Intel Labs China
{wangyk17@mails.,fcsun@}tsinghua.edu.cn, {duo.li,anbang.yao}@intel.com

Abstract. We propose a general method to train a single convolutional neural network which is capable of switching image resolutions at inference. Thus the running speed can be selected to meet various computational resource limits. Networks trained with the proposed method are named Resolution Switchable Networks (RS-Nets). The basic training framework shares network parameters for handling images which differ in resolution, yet keeps separate batch normalization layers. Though it is parameter-efficient in design, it leads to inconsistent accuracy variations at different resolutions, for which we provide a detailed analysis from the aspect of the train-test recognition discrepancy. A multi-resolution ensemble distillation is further designed, where a teacher is learnt on the fly as a weighted ensemble over resolutions. Thanks to the ensemble and knowledge distillation, RS-Nets enjoy accuracy improvements at a wide range of resolutions compared with individually trained models. Extensive experiments on the ImageNet dataset are provided, and we additionally consider quantization problems. Code and models are available at <https://github.com/yikaiw/RS-Nets>.

Keywords: Efficient Design; Multi-Resolution; Ensemble Distillation

1 Introduction

Convolutional Neural Networks (CNNs) have achieved great success on image recognition tasks [5,14], and well-trained recognition models usually need to be deployed on mobile phones, robots or autonomous vehicles [1,10]. To fit the resource constraints of devices, extensive research efforts have been devoted to balancing between accuracy and efficiency, by reducing computational complexities of models. Some of these methods adjust the structural configurations of networks, e.g., by adjusting the network depths [5], widths [10,31] or the convolutional blocks [18,33]. Besides that, adjusting the image resolution is another widely-used method for the accuracy-efficiency trade-off [9,10,16,24]. If input

* This work was done when Yikai Wang was an intern at Intel Labs China, supervised by Anbang Yao who is responsible for correspondence.

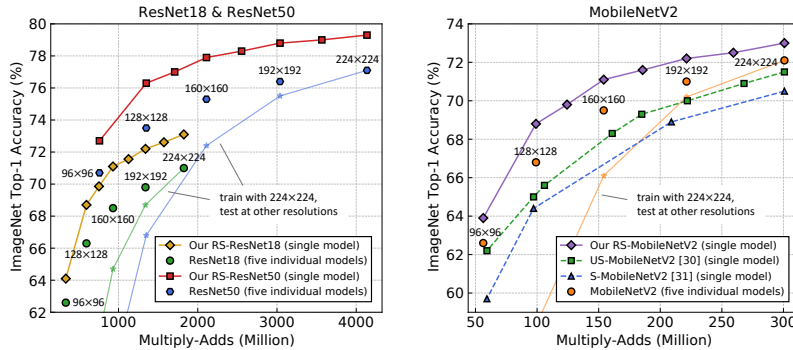


Fig. 1. ImageNet accuracy vs. FLOPs (Multiply-Adds) of our **single** models and the corresponding **sets** of individual models. A single RS-Net model is executable at each of the resolutions, and even achieves significantly higher accuracies than individual models. The results of two state-of-the-art switchable networks (switch by varying network widths) S-MobileNetV2 [31] and US-MobileNetV2 [30] are provided for comparison. Details are in Table 2 and Table 4.

images are downsized, all feature resolutions at different convolutional layers are reduced subsequently with the same ratio, and the computational cost of a model is nearly proportional to the image resolution ($H \times W$) [10]. However, for a common image recognition model, when the test image resolution differs from the resolution used for training, the accuracy quickly deteriorates [29]. To address this issue, existing works [9,24,28] train an individual model for each resolution. As a result, the total number of models to be trained and saved is proportional to the amount of resolutions considered at runtime. Besides the high storage costs, each time adjusting the resolution is accompanied with the additional latency to load another model which is trained with the target resolution.

The ability to switch the image resolution at inference meets a common need for real-life model deployments. By switching resolutions, the running speeds and costs are adjustable to flexibly handle the real-time latency and power requirements for different application scenarios or workloads. Besides, the flexible latency compatibility allows such model to be deployed on a wide range of resource-constrained platforms, which is friendly for application developers. In this paper, we focus on switching input resolutions for an image recognition model, and propose a general and economic method to improve overall accuracies. Models trained with our method are called **Resolution Switchable Networks (RS-Nets)**. Our contribution is composed of three parts.

First, we propose a parallel training framework where images with different resolutions are trained within a single model. As the resolution difference usually leads to the difference of activation statistics in a network [29], we adopt shared network parameters but privatized Batch Normalization layers (BNs) [12] for each resolution. Switching BNs enables the model to flexibly switch image resolutions, without needing to adjust other network parameters.

Second, we associate the multi-resolution interaction effects with a kind of train-test discrepancy (details in Section 3.2). Both our analysis and empirical

results reach an interesting conclusion that the parallel training framework tends to enlarge the accuracy gaps over different resolutions. On the one hand, accuracy promotions at high resolutions make a stronger teacher potentially available. On the other hand, the accuracy drop at the lower resolution indicates that the benefits of parallel training itself are limited. Both reasons encourage us to further propose a design of ensemble distillation to improve overall performance.

Third, to the best of our knowledge, we are the first to propose that a data-driven ensemble distillation can be learnt on the fly for image recognition, based on the same image instances with different resolutions. Regarding the supervised image recognition, the structure of our design is also different from existing ensemble or knowledge distillation works, as they focus on the knowledge transfer among different models, e.g., by stacking multiple models [34], multiple branches [15], pre-training a teacher model, or splitting the model into sub-models [30], while our model is single and shared, with little extra parameters.

Extensive experiments on the ImageNet dataset validate that RS-Nets are executable given different image resolutions at runtime, and achieve significant accuracy improvements at a wide range of resolutions compared with individually trained models. Illustrative results are provided in Fig. 1, which also verify that our proposed method can be generally applied to modern recognition backbones.

2 Related Work

Image Recognition. Image recognition acts as a benchmark to evaluate models and is a core task for computer vision. Advances on the large-scale image recognition datasets, like ImageNet [3], can translate to improved results on a number of other applications [4,13,20]. In order to enhance the model generalization for image recognition, data augmentation strategies, e.g., random-size crop, are adopted during training [5,6,26]. Besides, common models are usually trained and tested with fixed-resolution inputs. [29] shows that for a model trained with the default 224×224 resolution and tested at lower resolutions, the accuracy quickly deteriorates (e.g., drops 11.6% at test resolution 128×128 on ResNet50).

Accuracy-Efficiency Trade-Off. There have been many attempts to balance accuracy and efficiency by model scaling. Some of them adjust the structural configurations of networks. For example, ResNets [5] provide several choices of network depths from shallower to deep. MobileNets [10,24] and ShuffleNets [33] can reduce network widths by using smaller width multipliers. While some other works [9,10,16,24] reduce the computational complexity by decreasing image resolutions at input, which is also our focus. Modifying the resolution usually does not make changes to the number of network parameters, but significantly affects the computational complexity [10].

Knowledge Distillation. A student network can be improved by imitating feature representations or soft targets of a larger teacher network [7,17,34]. The teacher is usually pre-trained beforehand and fixed, and the knowledge is transferred in one direction [22]. Yet [34] introduces a two-way transfer between two




Fig. 2. Overall framework of training a RS-Net. Images with different resolutions are trained in parallel with shared Conv/FC layers and private BNs. The ensemble logit (z_0) is learnt on the fly as a weighted mean of logits (z_1, z_2, \dots, z_S), shown as green arrows. Knowledge distillations are shown as red arrows. For inference, one of the S forward paths is selected (according to the image resolution), with its corresponding BNs, for obtaining its corresponding prediction $p_s, s \in \{1, 2, \dots, S\}$. The ensemble and knowledge distillation are not needed during inference.

peer models. [25] performs mutual learning within one single network assisted by intermediate classifiers. [15] learns a native ensemble design based on multiple models for distillation. [30] conducts the knowledge distillation between the whole model and each split smaller model. Regarding the supervised image recognition, existing distillation works rely on different models, usually needing another teacher network with higher-capacity than low-capacity students. While our design is applied in a shared model, which is data-driven, collecting complementary knowledge from the same image instances with different resolutions.

3 Proposed Method

A schematic overview of our proposed method is shown in Fig. 2. In this section, we detail the insights and formulations of parallel training, interaction effects and ensemble distillation based on the multi-resolution setting.

3.1 Multi-Resolution Parallel Training

To make the description self-contained, we begin with the basic training of a CNN model. Given training samples, we crop and resize each sample to a fixed-resolution image \mathbf{x}^i . We denote network inputs as $\{(\mathbf{x}^i, y^i) | i \in \{1, 2, \dots, N\}\}$, where y^i is the ground truth which belongs to one of the C classes, and N is the amount of samples. Given network configurations with parameters θ , the predicted probability of the class c is denoted as $p(c|\mathbf{x}^i, \theta)$. The model is optimized with a cross-entropy loss defined as:

$$\mathcal{H}(\mathbf{x}, \mathbf{y}) = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C \delta(c, y^i) \log(p(c|\mathbf{x}^i, \theta)), \quad (1)$$

where $\delta(c, y^i)$ equals to 1 when $c = y^i$, otherwise 0.

In this part, we propose **multi-resolution parallel training**, or called **parallel training** for brevity, to train a single model which can switch image resolutions at runtime. During training, each image sample is randomly cropped and resized to several duplicate images with different resolutions. Suppose that there are S resolutions in total, the inputs can be written as $\{(\mathbf{x}_1^i, \mathbf{x}_2^i, \dots, \mathbf{x}_S^i, y^i) | i \in \{1, 2, \dots, N\}\}$. Recent CNNs for image recognition follow similar structures that all stack Convolutional (Conv) layers, a Global Average Pooling (GAP) layer and a Fully-Connected (FC) layer. In CNNs, if input images have different resolutions, the corresponding feature maps in all Conv layers will also vary in resolution. Thanks to GAP, features are transformed to a unified spatial dimension (1×1) with equal amount of channels, making it possible to be followed by a same FC layer. During our parallel training, we share parameters of Conv layers and the FC layer, and therefore the training for multiple resolutions can be realized in a single network. The loss function for parallel training is calculated as a summation of the cross-entropy losses:

$$\mathcal{L}_{cls} = \sum_{s=1}^S \mathcal{H}(\mathbf{x}_s, \mathbf{y}). \quad (2)$$

Specializing Batch Normalization layers (BNs) [12] is proved to be effective for efficient model adaption [2,19,31,35]. In image recognition tasks, resizing image results in different activation statistics in a network [29], including means and variances used in BNs. Thus during parallel training, we privatize BNs for each resolution. Results in the left panel of Fig. 4 verify the necessity of privatizing BNs. For the s^{th} resolution, each corresponding BN layer normalizes the channel-wise feature as follows:

$$\mathbf{y}'_s = \gamma_s \frac{\mathbf{y}_s - \boldsymbol{\mu}_s}{\sqrt{\boldsymbol{\sigma}_s^2 + \epsilon}} + \boldsymbol{\beta}_s, s \in \{1, 2, \dots, S\}, \quad (3)$$

where $\boldsymbol{\mu}_s$ and $\boldsymbol{\sigma}_s^2$ are running mean and variance; γ_s and $\boldsymbol{\beta}_s$ are learnable scale and bias. Switching these parameters enables the model to switch resolutions.

3.2 Multi-Resolution Interaction Effects

In this section, restricted to large-scale image datasets with fine-resolution images, we analyze the interaction effects of different resolutions under the parallel training framework. We start by posing a question: compared with individually trained models, how does parallel training affect test accuracies at different resolutions? As multi-resolution can be seen as a kind of data augmentation, we analyze from two aspects as follows.

The first aspect is straightforward. The model meets a wide range of image resolutions, which improves the generalization and reduces over-fitting. Thus if the setting of resolutions is suitable (e.g., not too diverse), the parallel training tends to bring overall accuracy gains at testing, especially for a high-capacity network such as ResNet50.

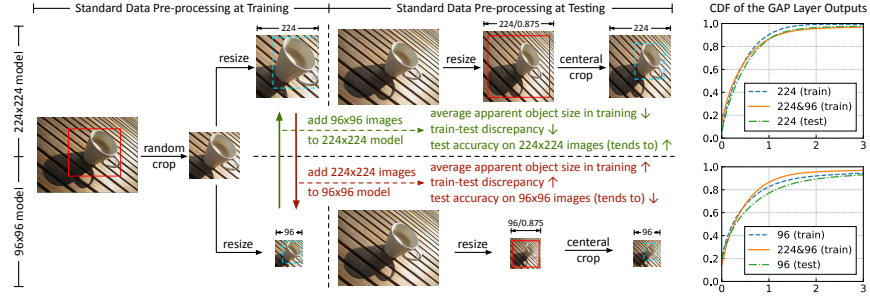


Fig. 3. Left: An illustration of interaction effects for the parallel training with two resolutions. Each red box indicates the region to be cropped, and the size of each blue dotted box is the apparent object (in this sample is a cup) size. For this example, in either one of the models, the apparent object size at testing is smaller than at training. [29] reveals that this relation still holds when averaging all data, which is called the train-test discrepancy. The data pre-processing for training or testing follows the standard image recognition method, which will be described in Section 4.1. **Right:** CDF curves for comparing the value distributions of feature activations. All curves are plotted on the validation dataset of ImageNet, but are based on different data pre-processing methods as annotated by (train) or (test).

The second aspect is based on the specialty of large-scale image recognition tasks, where objects of interest randomly occupy different portions of image areas and thus the random-size crop augmentation is used during training. [29] reveals that for common recognition works, as the random-size crop is used for training but not for testing, there exists a train-test **discrepancy** that the average “apparent object size” at testing is smaller than that at training. Besides, [29] achieves accuracy improvements by alleviating such discrepancy, but is accompanied with the costs of test-time resolution augmentations and finetuning. Note that we do not aim to modify the data pre-processing method to alleviate the discrepancy. Instead, we are inspired to use the concept of the discrepancy to analyze multi-resolution interaction effects. We take the parallel training with two resolutions 224×224 and 96×96 for example. According to the left panel of Fig. 3 (the analysis in colors), compared with the model using only 224×224 images, the parallel training can be seen as augmenting 96×96 images, which reduces the average apparent object size at training and thus alleviates the discrepancy. On the contrary, compared with the individual model using 96×96 images, augmenting 224×224 images increases the discrepancy. Thus in this aspect, this parallel training tends to increase the test accuracy at 224×224 (actually +1.6% for ResNet18, as shown in Fig. 5), while tends to reduce the test accuracy at 96×96 (-0.8%). The right panel of Fig. 3 plots the Cumulative Distribution Function (CDF)³ of output components of the Global Average Pooling (GAP) layer for a well-trained ResNet18 (as a ReLU layer is before the GAP, all components are nonnegative). We plot CDF to compare the value distributions of feature activations when using training or testing data pre-processing method.

³ The CDF of a random variable X is defined as $F_X(x) = P(X \leq x)$, for all $x \in \mathbb{R}$.

Table 1. Proportions (%) of validation samples that are correctly classified (\checkmark) at a resolution but are wrongly classified (\times) at another, based on ResNet18. All models are well-trained. Lower numbers correspond to better performance, and numbers which are larger than the base are colored red, otherwise green. Benefited from MRED, all the proportions in the last column decrease.

$\times \backslash \checkmark$	Individual Training (base)					Parallel Training					Parallel Training + MRED				
	224	192	160	128	96	224	192	160	128	96	224	192	160	128	96
224	-	5.7	5.5	5.3	4.7	-	3.0 ↓	3.4 ↓	3.8 ↓	3.9 ↓	-	2.5 ↓	2.9 ↓	3.4 ↓	3.6 ↓
192	6.9	-	6.0	5.5	5.2	4.1 ↓	-	3.3 ↓	3.7 ↓	3.8 ↓	3.4 ↓	-	2.8 ↓	3.3 ↓	3.4 ↓
160	8.1	7.2	-	5.9	5.4	6.0 ↓	4.7 ↓	-	3.5 ↓	3.6 ↓	5.1 ↓	3.9 ↓	-	2.9 ↓	3.2 ↓
128	9.8	8.9	8.1	-	5.8	9.3 ↓	8.2 ↓	6.6 ↓	-	3.5 ↓	7.3 ↓	6.4 ↓	5.0 ↓	-	3.1 ↓
96	13.2	12.5	11.4	9.7	-	15.2 ↑	14.2 ↑	12.7 ↑	9.8 ↑	-	12.2 ↓	11.4 ↓	10.2 ↓	7.9 ↓	-

The parallel training seems to narrow the train-test gap for the 224×224 model, but widen the gap for the 96×96 model.

We take ResNet18 as an example and summarize the two aforementioned aspects. For the parallel training with two resolutions, the test accuracy at the high resolution increases compared with its individual model, as the two aspects reach an agreement to a large degree. As for the lower resolution, we find that the accuracy slightly increases if the two resolutions are close, otherwise decreases. Similarly, when multiple resolutions are used for parallel training, test accuracies increase at high resolutions but may decrease at lower resolutions, compared with individual models. Results in Table 2 show that for the parallel training with five resolutions, accuracies only decrease at 96×96 but increase at the other four resolutions. Detailed results in Fig. 5 also verify our analysis.

For image recognition, although testing at a high resolution already tends to achieve a good accuracy, using the parallel training makes it even better. This finding opens up a possibility that a stronger teacher may be available in this framework, and seeking a design based on such teacher could be highly effective.

3.3 Multi-Resolution Ensemble Distillation

In this section, we propose a new design of ensemble distillation. Regarding the supervised image recognition, unlike conventional distillation works that rely on transferring knowledge among different models, ours is data-driven and can be applied in a shared model. Specifically, our design is learnt on the fly and the distillation is based on the same image instances with different resolutions.

As is commonly known, for image recognition tasks, models given a high resolution image are easy to capture fine-grained patterns, and thus achieve good performance [27,28]. However, according to the sample statistics in the middle column of Table 1, we find that there always exists a proportion of samples which are correctly classified at a low resolution but wrongly classified at another higher resolution. Such results indicate that model predictions at different image resolutions are complementary, and not always the higher resolution is better for each image sample. Therefore, we propose to learn a teacher on the fly as an ensemble of the predictions w.r.t. all resolutions, and conduct knowledge distillation to improve the overall performance. Our design is called **Multi-Resolution Ensemble Distillation (MRED)**.

During the training process of image recognition, for each input image \mathbf{x}^i , the probability of the class c is calculated using a softmax function:

$$p(c|\mathbf{x}^i, \boldsymbol{\theta}) = p(c|\mathbf{z}^i) = \frac{\exp(z_c^i)}{\sum_{j=1}^C \exp(z_j^i)}, c \in \{1, 2, \dots, C\}, \quad (4)$$

where \mathbf{z}^i is the logit, the unnormalized log probability outputted by the network, and probabilities over all classes can be denoted as the model prediction \mathbf{p} .

In the parallel training framework, each image is randomly cropped and resized to S images with different resolutions. To better benefit from MRED, these S images need to be resized from a same random crop, as illustrated in the left-most part of Fig. 2. The necessity will be verified in Section 4.3.

As each image sample is resized to S resolutions, there are S corresponding logits $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_S$. We learn a group of importance scores $\boldsymbol{\alpha} = [\alpha_1 \alpha_2 \dots \alpha_S]$, satisfying $\boldsymbol{\alpha} \geq 0$, $\sum_{s=1}^S \alpha_s = 1$, which can be easily implemented with a softmax function. We then calculate an ensemble logit \mathbf{z}_0 as the weighted summation of the S logits:

$$\mathbf{z}_0 = \sum_{s=1}^S \alpha_s \mathbf{z}_s. \quad (5)$$

To optimize $\boldsymbol{\alpha}$, we temporally froze the gradients of the logits $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_S$. Based on the ensemble logit \mathbf{z}_0 , the corresponding prediction \mathbf{p}_0 , called ensemble prediction, can be calculated via Eq. 4. Then $\boldsymbol{\alpha}$ is optimized using a cross-entropy loss between \mathbf{p}_0 and the ground truth, which we call the ensemble loss \mathcal{L}_{ens} :

$$\mathcal{L}_{ens} = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C \delta(c, y^i) \log(p(c|\mathbf{z}_0^i)). \quad (6)$$

In knowledge distillation works, to quantify the alignment between a teacher prediction \mathbf{p}_t and a student prediction \mathbf{p}_s , Kullback Leibler (KL) divergence is usually used:

$$\mathcal{D}_{kl}(\mathbf{p}_t \parallel \mathbf{p}_s) = \frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C p(c|\mathbf{z}_t^i) \log \frac{p(c|\mathbf{z}_t^i)}{p(c|\mathbf{z}_s^i)}. \quad (7)$$

We force predicationations at different resolutions to mimic the learnt ensemble prediction \mathbf{p}_0 , and thus the distillation loss \mathcal{L}_{dis} could be obtained as:

$$\mathcal{L}_{dis} = \sum_{s=1}^S \mathcal{D}_{kl}(\mathbf{p}_0 \parallel \mathbf{p}_s). \quad (8)$$

Finally, the overall loss function is a summation of the classification loss, the ensemble loss and the distillation loss, without needing to tune any extra weighted parameters:

$$\mathcal{L} = \mathcal{L}_{cls} + \mathcal{L}_{ens} + \mathcal{L}_{dis}, \quad (9)$$

where in practical, optimizing \mathcal{L}_{ens} only updates $\boldsymbol{\alpha}$, with all network weights temporally frozen; optimizing \mathcal{L}_{cls} and \mathcal{L}_{dis} updates network weights.

We denote the method with Eq. 8 as our vanilla-version MRED. Under the parallel training framework, as the accuracy at a high resolution is usually better than at a lower resolution, accuracies can be further improved by offering dense guidance from predications at high resolutions toward predictions at lower resolutions. Thus the distillation loss can be extended to be a generalized one:

$$\mathcal{L}_{dis} = \frac{2}{S+1} \sum_{t=0}^{S-1} \sum_{s=t+1}^S \mathcal{D}_{kl}(\mathbf{p}_t \| \mathbf{p}_s), \quad (10)$$

where the index t starts from 0 referring to the ensemble term; as the summation results in $S(S+1)/2$ components in total, we multiply \mathcal{L}_{dis} by a constant ratio $2/(S+1)$ to keep its range the same as \mathcal{L}_{cls} . We denote the method with Eq. 10 as our full-version MRED, which is used in our experiments by default.

The proposed design involves negligible extra parameters (only S scalars), without needing extra models. Models trained with the parallel training framework and MRED are named Resolution Switchable Networks (RS-Nets). An overall framework of training a RS-Net is illustrated in Fig. 2. **During inference**, the network only performs one forward calculation at a given resolution, without ensemble or distillation, and thus both the computational complexity and the amount of parameters equal to a conventional image recognition model.

4 Experiments

We perform experiments on ImageNet (ILSVRC12) [3,23], a widely-used image recognition dataset containing about 1.2 million training images and 50 thousand validation images, where each image is annotated as one of 1000 categories. Experiments are conducted with prevailing CNN architectures including a lightweight model MobileNetV2 [24] and ResNets [5], where a basic-block model ResNet18 and a bottleneck-block model ResNet50 are both considered. Besides, we also evaluate our method in handling network quantization problems, where we consider different kinds of bit-widths.

4.1 Implementation Details

Our basic experiments are implemented with PyTorch [21]. For quantization tasks, we apply our method to LQ-Nets [32] which show state-of-the-art performance in training CNNs with low-precision weights or both weights and activations.

We set $\mathbb{S} = \{224 \times 224, 192 \times 192, 160 \times 160, 128 \times 128, 96 \times 96\}$, as commonly adopted in a number of existing works [10,16,24]. During training, we pre-process the data for augmentation with an area ratio (cropped area/original area) uniformly sampled in $[0.08, 1.0]$, an aspect ratio $[3/4, 4/3]$ and a horizontal flipping. We resize images with the bilinear interpolation. Note that both $[0.08, 1.0]$ and $[3/4, 4/3]$ follow the standard data augmentation strategies for ImageNet [11,26,29], e.g., *RandomResizedCrop* in PyTorch uses such setting as default. During validation, we first resize images with the bilinear interpolation to every resolution in \mathbb{S} divided by 0.875 [8,16], and then feed central regions to models.

Table 2. Basic results comparison on ImageNet. We report top-1/top-5 accuracies (%), top-1 accuracy gains (%) over individual models (I-Nets), Multiply-Adds (MAdds) and total parameters (params). All experiments use the same data pre-processing methods. Our baseline results are slightly higher than the original papers [5,24].

Network	Resolution	MAdds	I-Nets (base)		I-224	Our Parallel	Our RS-Net
ResNet18	224 × 224	1.82G	71.0 / 90.0	71.0 / 90.0		73.0 / 90.9 (+2.0)	73.1 / 91.0 (+2.1)
	192 × 192	1.34G	69.8 / 89.4	68.7 / 88.5 (-1.1)		71.7 / 90.3 (+1.9)	72.2 / 90.6 (+2.4)
	160 × 160	931M	68.5 / 88.2	64.7 / 85.9 (-5.2)		70.4 / 89.6 (+1.9)	71.1 / 90.1 (+2.6)
	128 × 128	596M	66.3 / 86.8	56.8 / 80.0 (-9.5)		67.5 / 87.8 (+1.2)	68.7 / 88.5 (+2.4)
	96 × 96	335M	62.6 / 84.1	42.5 / 67.9 (-20.1)		61.5 / 83.5 (-1.1)	64.1 / 85.3 (+1.5)
	Total Params		55.74M	11.15M		11.18M	11.18M
ResNet50	224 × 224	4.14G	77.1 / 93.4	77.1 / 93.4		78.9 / 94.4 (+1.8)	79.3 / 94.6 (+2.2)
	192 × 192	3.04G	76.4 / 93.2	75.5 / 92.5 (-0.9)		78.1 / 94.0 (+1.7)	78.8 / 94.4 (+2.4)
	160 × 160	2.11G	75.3 / 92.4	72.4 / 90.7 (-2.9)		76.9 / 93.1 (+1.6)	77.9 / 93.9 (+2.6)
	128 × 128	1.35G	73.5 / 91.4	66.8 / 87.0 (-6.7)		74.9 / 92.1 (+1.4)	76.3 / 93.0 (+2.8)
	96 × 96	760M	70.7 / 89.8	54.9 / 78.2 (-15.8)		70.2 / 89.4 (-0.5)	72.7 / 91.0 (+2.0)
	Total Params		121.87M	24.37M		24.58M	24.58M
M-NetV2	224 × 224	301M	72.1 / 90.5	72.1 / 90.5		72.8 / 90.9 (+0.7)	73.0 / 90.8 (+0.9)
	192 × 192	221M	71.0 / 89.8	70.2 / 89.1 (-0.9)		71.7 / 90.2 (+0.7)	72.2 / 90.5 (+1.2)
	160 × 160	154M	69.5 / 88.9	66.1 / 86.3 (-3.2)		70.1 / 89.2 (+0.6)	71.1 / 90.2 (+1.6)
	128 × 128	99M	66.8 / 87.0	58.3 / 81.2 (-8.5)		67.3 / 87.2 (+0.5)	68.8 / 88.2 (+2.0)
	96 × 96	56M	62.6 / 84.0	43.9 / 69.1 (-18.7)		61.4 / 83.3 (-1.2)	63.9 / 84.9 (+1.3)
	Total Params		16.71M	3.34M		3.47M	3.47M

Networks are trained from scratch with random initializations. We set the batch size to 256, and use a SGD optimizer with a momentum 0.9. For standard ResNets, we train 120 epochs and the learning rate is annealed from 0.1 to 0 with a cosine scheduling [6]. For MobileNetV2, we train 150 epochs and the learning rate is annealed from 0.05 to 0 with a cosine scheduling. For quantized ResNets, we follow the settings in LQ-Nets [32], which train 120 epochs and the learning rate is initialized to 0.1 and divided by 10 at 30, 60, 85, 95, 105 epochs. The weight decay rate is set to 1e-4 for all ResNets and 4e-5 for MobileNetV2.

4.2 Results

As mentioned in Section 1, common works with multi-resolution settings train and deploy multiple individual models separately for different resolutions. We denote these individual models as **I-Nets**, which are set as baselines. We use I-{resolution} to represent each individual model, e.g., I-224.

Basic Results. In Table 2, we report results on ResNet18, ResNet50 and MobileNetV2 (M-NetV2 for short). Besides I-Nets, we also report accuracies at five resolutions using the individual model which is trained with the largest resolution (I-224). For our proposed method, we provide separate results of the parallel training (parallel) and the overall design (RS-Net). As mentioned in Section 1, I-Nets need several times of parameter amount and high latencies for switching across models. We also cannot rely on an individual model to switch the image resolutions, as accuracies of I-224 are much lower than I-Nets at other resolutions (e.g., 15%~20% accuracy drop at the resolution 96 × 96). Similarly, each of the other individual models also suffers from serious accuracy drops, as can

Table 3. Results comparison for quantization tasks. We report top-1/top-5 accuracies (%) and top-1 accuracy gains (%) over individual models (I-Nets). All experiments are performed under the same training settings following LQ-Nets.

Network	Resolution	Bit-width (W/A): 2 / 32		Bit-width (W/A): 2 / 2	
		I-Nets (base)	Our RS-Net	I-Nets (base)	Our RS-Net
Quantized ResNet18	224 × 224	68.0 / 88.0	68.8 / 88.4 (+0.8)	64.9 / 86.0	65.8 / 86.4 (+0.9)
	192 × 192	66.4 / 86.9	67.6 / 87.8 (+1.2)	63.1 / 84.7	64.8 / 85.8 (+1.7)
	160 × 160	64.5 / 85.5	66.0 / 86.5 (+1.5)	61.1 / 83.3	62.9 / 84.2 (+1.8)
	128 × 128	61.5 / 83.4	63.1 / 84.5 (+1.6)	58.1 / 80.8	59.3 / 81.9 (+1.2)
	96 × 96	56.3 / 79.4	56.6 / 79.9 (+0.3)	52.3 / 76.4	52.5 / 76.7 (+0.2)
Quantized ResNet50	224 × 224	74.6 / 92.2	76.0 / 92.8 (+1.4)	72.2 / 90.8	74.0 / 91.5 (+1.8)
	192 × 192	73.5 / 91.3	75.1 / 92.4 (+1.6)	70.9 / 89.8	73.1 / 91.0 (+2.2)
	160 × 160	71.9 / 90.4	73.8 / 91.6 (+1.9)	69.0 / 88.5	71.4 / 90.0 (+2.4)
	128 × 128	69.6 / 88.9	71.7 / 90.2 (+2.1)	66.6 / 86.9	68.9 / 88.3 (+2.3)
	96 × 96	65.5 / 86.0	67.3 / 87.4 (+1.8)	61.7 / 83.4	63.4 / 84.7 (+1.7)

Table 4. Top-1 accuracies (%) on MobileNetV2. Individual models (I-Nets-w, adjust via network width; I-Nets-r, adjust via resolution) and our RS-Net are trained under the same settings. Results of S-MobileNetV2 (S) [31] and US-MobileNetV2 (US) [30] are from the original papers.

Width	MAdds	I-Nets-w	S [31]	US [30]	Resolution	MAdds	I-Nets-r	Ours
1.0×	301M	72.1	70.5	71.5	224 × 224	301M	72.1	73.0
0.8×	222M	69.8	-	70.0	192 × 192	221M	71.0	72.2
0.65×	161M	68.0	-	68.3	160 × 160	154M	69.5	71.1
0.5×	97M	64.6	64.4	65.0	128 × 128	99M	66.8	68.8
0.35×	59M	60.1	59.7	62.2	96 × 96	56M	62.6	63.9
Model Size		5×	1×	1×	Model Size		5×	1×

be seen in the right panel of Fig. 4. Our parallel training brings accuracy improvements at the four larger resolutions, while accuracies at 96×96 decrease, and the reason is previously analyzed in Section 3.2. Compared with I-Nets, our RS-Net achieves large improvements at all resolutions with only 1/5 parameters. For example, the RS-Net with ResNet50 obtains about 2.4% absolute top-1 accuracy gains on average across five resolutions. Note that the number of FLOPs (Multiply-Adds) is nearly proportional to the image resolution [10]. Regarding ResNet18 and ResNet50, accuracies at 160×160 of our RS-Nets even surpass the accuracies of I-Nets at 224×224 , significantly reducing about 49% FLOPs at runtime. Similarly, for MobileNetV2, the accuracy at 192×192 of RS-Net surpasses the accuracy of I-Nets at 224×224 , reducing about 26% FLOPs.

Quantization. We further explore the generalization of our method to more challenging quantization problems, and we apply our method to LQ-Nets [32]. Experiments are performed under two typical kinds of quantization settings, including the quantization on weights (2/32) and the more extremely compressed quantization on both weights and activations (2/2). Results of I-Nets and each RS-Net based on quantized ResNets are reported in Table 3. Again, each RS-Net outperforms the corresponding I-Nets at all resolutions. For quantization problems, as I-Nets cannot force the quantized parameter values of each indi-

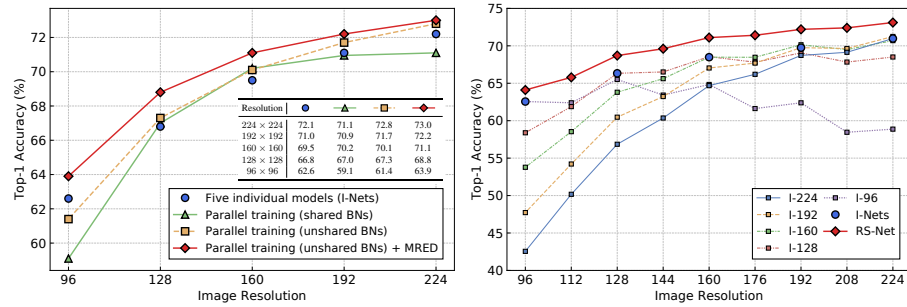


Fig. 4. Left: Comparison of parallel trainings with shared BNs and unshared (private) BNs, based on MobileNetV2. Individual models and Parallel (unshared BNs) + MRED (i.e., RS-Net) are provided for reference. **Right:** Comparison of our RS-Net and each individual model (from I-224 to I-96) tested at denser resolutions (the interval is 16), based on ResNet18. Each individual model suffers from serious accuracy drops at other resolutions, but our RS-Net avoids this issue.

vidual model to be the same, 2-bit weights in I-Nets are practically stored with more digits than a 4-bit model⁴, while the RS-Net avoids such issue. For quantization, accuracy gains of RS-Net to ResNet50 are more obvious than those to ResNet18, we conjecture that under compressed conditions, ResNet50 better bridges the network capacity and the augmented data resolutions.

Switchable Models Comparison. Based on MobileNetV2, results in Table 4 indicate that under comparable FLOPs (Multiply-Adds), adjusting image resolutions (see I-Nets-r) achieves higher accuracies than adjusting network widths (see I-Nets-w). For example, adjusting resolution to 96×96 brings 2.5% higher absolute accuracy than adjusting width to $0.35\times$, with even lower FLOPs. Results of S-MobileNetV2 (S) [31] and US-MobileNetV2 (US) [30] are also provided for comparison. As we can see, our RS-Net significantly outperforms S and US at all given FLOPs, achieving 1.5%~4.4% absolute gains. Although both S and US can also adjust the accuracy-efficiency trade-off, they have marginal gains or even accuracy drops (e.g., at width $1.0\times$) compared with their baseline I-Nets-w. Our model shows large gains compared with our baseline I-Nets-r (even they are mostly stronger than I-Nets-w). Note that adjusting resolutions does not conflict with adjusting widths, and both methods can be potentially combined together.

4.3 Ablation Study

Importance of Using Private BNs. A quick question is that, why not share BNs as well? Based on MobileNetV2, the left panel of Fig. 4 shows that during the parallel training, privatizing BNs achieves higher accuracies than sharing BNs, especially at both the highest resolution (+1.7%) and the lowest (+2.3%). When BNs are shared, activation statistics of different image resolutions are averaged, which differ from the real statistics especially at two ends of resolutions.

⁴ The total bit number of five models with individual 2-bit weights is $\log_2(2^2 \times 5) \approx 4.3$.




Fig. 5. Absolute top-1 accuracy variations (%) (compared with individual models) of parallel trainings, based on ResNet18. The result of each individual model (from I-96 to I-224) is used as the baseline. We use single numbers to represent the image resolutions.

Tested at New Resolutions. One may concern that if a RS-Net can be tested at a new resolution. In the right panel of Fig. 4, we test models at different resolutions with a smaller interval, using ResNet18. We follow a simple method to handle a resolution which is not involved in training. Suppose the resolution is sandwiched between two of the five training resolutions which correspond to two groups of BN parameters, we apply a linear interpolation on these two groups and obtain a new group of parameters, which are used for the given resolution. We observe that the RS-Net maintains high accuracies. RS-Net suppresses the serious accuracy drops which exist in every individual model.

Multi-Resolution Interaction Effects. This part is for verifying the analysis in Section 3.2, and we do not apply MRED here. Parallel training results of ResNet18 are provided in Fig. 5, which show top-1 accuracy variations compared with each individual model. The left panel of Fig. 5 illustrates the parallel training with two resolutions, where all accuracies increase at 224×224 . As the gap of the two resolutions increases, the accuracy variation at the lower resolution decreases. For example, compared with I-224 and I-96 respectively, the parallel training with 224×224 and 96×96 images increases the accuracy at 224×224 from 71.0% to 72.6%, but decreases the accuracy at 96×96 from 62.6% to 61.8%. The right panel of Fig. 5 illustrates the parallel training with multiple resolutions, where we observe that most accuracies are improved and accuracies at the lowest resolutions may decrease. These results verify our analysis.

Verifying Each Element of Ensemble Distillation. In Section 3.3, we define two versions of MRED. The vanilla-version has only the distillation paths starting from the ensemble prediction p_0 toward all the other predictions, while the full-version has additional paths from predictions at high resolutions toward predictions at lower resolutions. In Table 5, we compare the performance of the two versions as well as two other variants that omit the distillations from p_0 , based on ResNet18. Results indicate that all kinds of the proposed distillations are indispensable. Besides, we also emphasize in Section 3.3 that for training a RS-Net, each image should be randomly cropped only once (called single-crop) and then resized to multiple resolutions (called multi-resolution). Results in Table 6 indicate that using multi-crop (applying a random crop individually for each resolution) will weaken the benefits of MRED, as accuracies

Table 5. Top-1 accuracies (%) of different kinds of distillations and their accuracy variations compared with the full-version MRED, based on ResNet18. Definitions of $\mathbf{p}_s, s \in \{1, 2, \dots, S\}$ are referred to Section 3.3. The prediction \mathbf{p}_1 corresponds to the largest resolution 224×224 .

Resolution	Full-version MRED (base)	Vanilla-version MRED	Only distillations from $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_{S-1}$	Only distillations from \mathbf{p}_1
224×224	73.1	73.0 (-0.1)	72.3 (-0.8)	72.2 (-0.9)
192×192	72.2	72.1 (-0.1)	71.2 (-1.0)	71.7 (-0.5)
160×160	71.1	70.8 (-0.3)	70.2 (-0.9)	70.5 (-0.6)
128×128	68.7	68.3 (-0.4)	68.1 (-0.6)	68.1 (-0.6)
96×96	64.1	63.7 (-0.4)	63.5 (-0.6)	63.1 (-1.0)

Table 6. Top-1 accuracies (%) comparison to verify the importance of using multi-resolution and single-crop, based on ResNet18. Multi-crop refers to applying the random crop individually for each resolution. Single-resolution refers to using five identical resolutions in a RS-Net instead of the original multi-resolution setting.

Resolution	Multi-resolution Single-crop (base)	Multi-resolution Multi-crop	Single-resolution Multi-crop	Single-resolution Single-crop
224×224	73.1	73.1 (-0.0)	72.1 (-1.0)	71.0 (-2.1)
192×192	72.2	72.0 (-0.2)	71.1 (-1.1)	69.8 (-2.4)
160×160	71.1	70.7 (-0.4)	69.6 (-1.5)	68.5 (-2.6)
128×128	68.7	68.0 (-0.7)	67.6 (-1.1)	66.3 (-2.4)
96×96	64.1	62.8 (-1.3)	63.3 (-0.8)	62.6 (-1.5)
Model Size	$1 \times$	$1 \times$	$5 \times$	$5 \times$

at low resolutions are lower compared with using single-crop. We also verify the importance of multi-resolution by replacing the multi-resolution setting by five identical resolutions, called single-resolution. As resolutions are identical, an individual experiment is needed for each resolution. Results indicate that applying ensemble distillation to predictions of different crops has very limited benefits compared with applying it to predictions of different resolutions.

5 Conclusions

We introduce a general method to train a CNN which can switch the resolution during inference, and thus its running speed can be selected to fit different computational constraints. Specifically, we propose the parallel training to handle multi-resolution images, using shared parameters and private BNs. We analyze the interaction effects of resolutions from the perspective of train-test discrepancy. And we propose to learn an ensemble distillation based on the same image instances with different resolutions, which improves accuracies to a large extent.

Acknowledgement

This work is jointly supported by the National Science Foundation of China (NSFC) and the German Research Foundation (DFG) in project Cross Modal Learning, NSFC 61621136008/DFG TRR-169. We thank Aojun Zhou for the insightful discussions.

References

1. Cai, H., Zhu, L., Han, S.: Proxylessnas: Direct neural architecture search on target task and hardware. In: ICLR (2019)
2. Chang, W., You, T., Seo, S., Kwak, S., Han, B.: Domain-specific batch normalization for unsupervised domain adaptation. In: CVPR (2019)
3. Deng, J., Dong, W., Socher, R., Li, L., Li, K., Li, F.: Imagenet: A large-scale hierarchical image database. In: CVPR (2009)
4. Gordo, A., Almazán, J., Revaud, J., Larlus, D.: End-to-end learning of deep visual representations for image retrieval. IJCV (2017)
5. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
6. He, T., Zhang, Z., Zhang, H., Zhang, Z., Xie, J., Li, M.: Bag of tricks for image classification with convolutional neural networks. In: CVPR (2019)
7. Hinton, G.E., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015)
8. Hoffer, E., Weinstein, B., Hubara, I., Ben-Nun, T., Hoefler, T., Soudry, D.: Mix & match: training convnets with mixed image sizes for improved accuracy, speed and scale resiliency. arXiv preprint arXiv:1908.08986 (2019)
9. Howard, A., Sandler, M., Chu, G., Chen, L., Chen, B., Tan, M., Wang, W., Zhu, Y., Pang, R., Vasudevan, V., Le, Q.V., Adam, H.: Searching for mobilenetv3. In: ICCV (2019)
10. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861 (2017)
11. Huang, G., Liu, Z., van der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: CVPR (2017)
12. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: ICML (2015)
13. Kornblith, S., Shlens, J., Le, Q.V.: Do better imagenet models transfer better? In: CVPR (2019)
14. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS (2012)
15. Lan, X., Zhu, X., Gong, S.: Knowledge distillation by on-the-fly native ensemble. In: NeurIPS (2018)
16. Li, D., Zhou, A., Yao, A.: Hbonet: Harmonious bottleneck on two orthogonal dimensions. In: ICCV (2019)
17. Li, Z., Hoiem, D.: Learning without forgetting. In: ECCV (2016)
18. Ma, N., Zhang, X., Zheng, H., Sun, J.: Shufflenet V2: practical guidelines for efficient CNN architecture design. In: ECCV (2018)
19. Mudrakarta, P.K., Sandler, M., Zhmoginov, A., Howard, A.G.: K for the price of 1: Parameter-efficient multi-task and transfer learning. In: ICLR (2019)
20. Oquab, M., Bottou, L., Laptev, I., Sivic, J.: Learning and transferring mid-level image representations using convolutional neural networks. In: CVPR (2014)
21. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: NeurIPS (2019)

22. Romero, A., Ballas, N., Kahou, S.E., Chassang, A., Gatta, C., Bengio, Y.: Fitnets: Hints for thin deep nets. In: ICLR (2015)
23. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M.S., Berg, A.C., Li, F.: Imagenet large scale visual recognition challenge. IJCV (2015)
24. Sandler, M., Howard, A.G., Zhu, M., Zhmoginov, A., Chen, L.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: CVPR (2018)
25. Sun, D., Yao, A., Zhou, A., Zhao, H.: Deeply-supervised knowledge synergy. In: CVPR (2019)
26. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S.E., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: CVPR (2015)
27. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: CVPR (2016)
28. Tan, M., Le, Q.V.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: ICML (2019)
29. Touvron, H., Vedaldi, A., Douze, M., Jégou, H.: Fixing the train-test resolution discrepancy. In: NeurIPS (2019)
30. Yu, J., Huang, T.S.: Universally slimmable networks and improved training techniques. In: ICCV (2019)
31. Yu, J., Yang, L., Xu, N., Yang, J., Huang, T.S.: Slimmable neural networks. In: ICLR (2019)
32. Zhang, D., Yang, J., Ye, D., Hua, G.: Lq-nets: Learned quantization for highly accurate and compact deep neural networks. In: ECCV (2018)
33. Zhang, X., Zhou, X., Lin, M., Sun, J.: Shufflenet: An extremely efficient convolutional neural network for mobile devices. In: CVPR (2018)
34. Zhang, Y., Xiang, T., Hospedales, T.M., Lu, H.: Deep mutual learning. In: CVPR (2018)
35. Zhou, A., Ma, Y., Li, Y., Zhang, X., Luo, P.: Towards improving generalization of deep networks via consistent normalization. arXiv preprint arXiv:1909.00182 (2019)