

Compact and Optimal Deep Learning with Recurrent Parameter Generators

Jiayun Wang*¹ Yubei Chen*² Stella X. Yu¹ Brian Cheung⁴ Yann LeCun^{2,3}

¹UC Berkeley / ICSI ²Facebook AI Research ³New York University ⁴MIT CSAIL & BCS
 {peterwg,stellayu}@berkeley.edu {yubeic,yann}@fb.com cheungb@mit.edu

Abstract

Deep learning has achieved tremendous success by training increasingly large models, which are then compressed for practical deployment. We propose a drastically different approach to compact and optimal deep learning: We decouple the Degrees of freedom (DoF) and the actual number of parameters of a model, optimize a small DoF with predefined random linear constraints for a large model of an arbitrary architecture, in one-stage end-to-end learning.

Specifically, we create a recurrent parameter generator (RPG), which repeatedly fetches parameters from a ring and unpacks them onto a large model with random permutation and sign flipping to promote parameter decorrelation. We show that gradient descent can automatically find the best model under constraints with in fact faster convergence.

Our extensive experimentation reveals a log-linear relationship between model DoF and accuracy. Our RPG demonstrates remarkable DoF reduction, and can be further pruned and quantized for additional run-time performance gain. For example, in terms of top-1 accuracy on ImageNet, RPG achieves 96% of ResNet18’s performance with only 18% DoF (the equivalent of one convolutional layer) and 52% of ResNet34’s performance with only 0.25% DoF! Our work shows significant potential of constrained neural optimization in compact and optimal deep learning.

1. Introduction

Deep neural networks as general optimization tools have achieved great success with increasingly more training data, deeper and larger neural networks: A recently developed NLP model, GPT-3 [8], has astonishing 175 billion parameters! While the model performance generally scales with the number of parameters [29], with parameters outnumbering training data, the model is significantly over-parameterized.

Many approaches have been proposed to remove redundancy in trained large models: neural network pruning

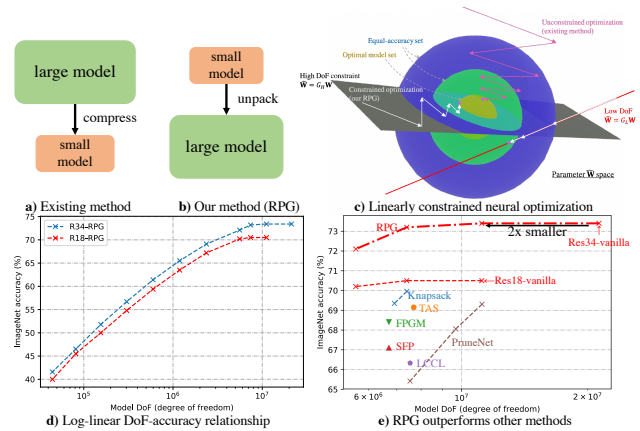


Figure 1: We propose a novel approach to compact and optimal deep learning by decoupling model DoF and model parameters. **a)** Existing methods first finds the optimal in a large model space and then compress it for practical deployment. **b)** We propose to start with a small (DoF) model of free parameters, use recurrent parameter generator (RPG) to unpack them onto a large model with predefined random linear projections. **c)** Gradient descent finds the optimal model of a small DoF under these linear constraints with faster converge than training the large unpacked model itself (Fig.5b). If the DoF is too small, the optimal large model may fall out of the constrained subspace. However, at a sufficiently large DoF, RPG gets rid of redundancy and often finds a model with little loss in accuracy. **d)** RPG reveals a log-linear relationship between model DoF and accuracy. **e)** RPG achieves the same ImageNet accuracy with half of the ResNet-vanilla DoF. RPG also outperforms other state-of-the-art compression approaches.

[39, 23, 42], efficient network design spaces [30, 33, 51], parameter regularization [59, 60, 52, 47], model quantization [31, 50, 43], neural architecture search [70, 10, 58], recurrent models [4, 5, 62], multi-task feature encoding [49, 24], etc. Pruning-based model compression dates back to the late 80s [45, 39] and has enjoyed recent resurgence [23, 7]. They remove unimportant parameters from a pre-trained model and can achieve significant model compression.

Our work is a departure from mainstream approaches towards model optimization and parameter reduction: rather than compressing a large model, we directly optimize a lean model with a small set of free parameters (number of free

*indicates equal contribution.

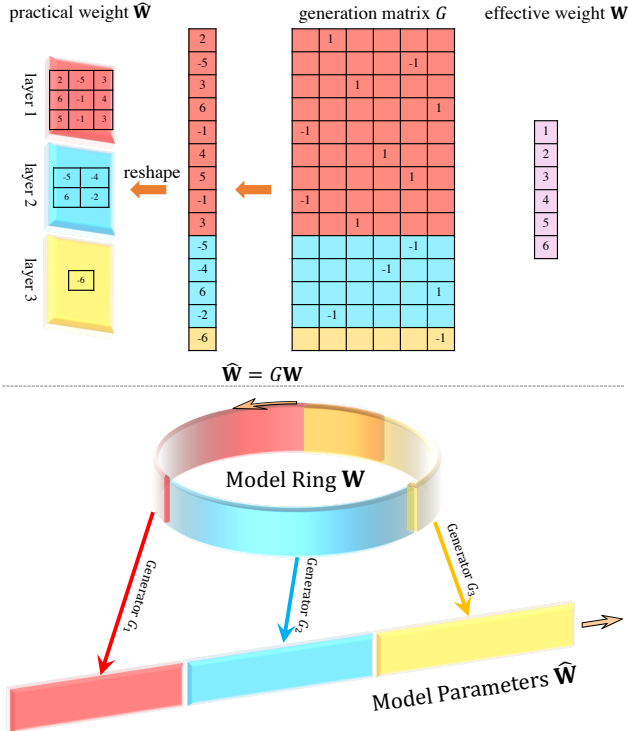


Figure 2: **Upper:** Networks are optimized with a linear constraint $\hat{\mathbf{W}} = \mathbf{G}\mathbf{W}$, where the constrained parameter $\hat{\mathbf{W}}$ of each network layer was generated by the generating matrix \mathbf{G} from the free parameter \mathbf{W} , which is directly optimized. $\hat{\mathbf{W}}$ is unpacked large model parameter while the size of \mathbf{W} is the model DoF. **Lower:** This paper discusses a specific format of parameter generation, recurrent parameter generator (RPG). RPG shares a fixed set of parameters in a ring and uses them to generate parameters of different parts of a neural network, whereas in the standard neural network, all the parameters are independent of each other, so the model gets bigger as it gets deeper. The **third** section of the model starts to overlap with the **first** section in the model ring, and all later layers share generating parameters for possibly multiple times.

parameters equal to degree of freedom of the model, or DoF), which can be linearly unpacked to a large model. Training the large model can be viewed as solving a neural optimization with a set of predefined linear constraints. One benefit of constrained neural optimization we observe is that it leads to a faster convergence rate (Section 5.6). Specifically, we define different layers in a neural network based on a fixed amount of DoF, which we call *recurrent parameter generator* (RPG). That is, we differentiate the number of model parameters and DoF. Traditionally, model parameters are treated independently of each other; the total number of parameters equals DoF. However, by tapping into how a core set of free parameters can be assigned to the neural network model, we can develop a large model of many parameters, which are linearly constrained by the small set of free parameters.

There is excess capacity in neural networks independent

of how and where the parameters are used in the network, even at the level of individual scalar values. Surprisingly, backpropagation training of a deep network is able to cope with that the same parameter can be assigned to multiple random locations in the network without significantly impacting model performance. Our extensive experiments show that a large neural network does not need to be over-parameterized to achieve competitive performance. Particularly, a ResNet18 can be implemented with DoF equivalent to one convolution layer in a ResNet18-vanilla ($4.72\times$ DoF reduction) and still achieves 67.2% ImageNet top-1 accuracy. The proposed method is also extremely flexible in reducing model DoF. In some sense, the proposed RPG method can be viewed as an automatic model DoF reduction technique, which explores the optimal accuracy-parameter trade-off. When we reduce the model DoF, RPG demonstrates graceful performance degradation, and its compression results are frequently on par with the SOTA pruning methods besides the flexibility. Even if we reduce the Res18 backbone DoF to 36K, which is about $300\times$ reduction, ResNet18 can still achieve 40% ImageNet top-1 accuracy. Further, we show RPG can be quantized and pruned to improve FLOPs and runtime with relatively mild accuracy drops.

To summarize, we make three contributions: 1. We provide a new perspective towards automatic model size reduction: we define a neural network with certain DoF with random linear constraints. We discover that gradient descent can automatically solve constrained optimization for the best model with a faster convergence rate. This constrained neural optimization perspective is likely to benefit many other applications. 2. We propose the recurrent parameter generator (RPG), which decouples the network architecture and the network DoF. We can flexibly choose any desired DoF to construct the network given a specific neural network architecture. 3. By separating network architectures from parameters, RPG becomes a tool to understand the relationship between the model DoF and the network performance. We observe an empirical log-linear DoF-Accuracy relationship.

2. Related Work

Many works study model DoF reduction or compression. We discuss each one and its relationship to our work.

Model Pruning, Neural Architecture Search, and Quantization. Model pruning seeks to remove unimportant parameters in a trained model. Recently, it's proposed to use neural architecture search as coarse-grained model pruning [68, 16]. Another related effort is network quantization [31, 50, 43], which seeks to reduce the bits used for each parameter and can frequently reduce the model size by $4\times$ with minimal accuracy drop. More recently, [14] presents a framework for analyzing model scaling strategies that consider network properties such as FLOPs and activations.

Parameter Regularization and Priors. Regularization has been widely used to reduce model redundancy [38, 47], al-




Figure 3: We demonstrate the effectiveness of RPG on various applications including image classification (**Left**), human pose estimation (**Middle**), and multitask regression (**Right**). RPGs are shared at multiple scales: a network can either have a global RPG or multiple local RPGs that are shared within blocks or sub-networks.

leviate overfitting [52, 59], and ensure desired mathematical regularity [60]. RPG can be viewed as a parameter regularization in the sense that weight sharing poses many equality constraints to weights and regularizes weights to a low-dimensional space. HyperNeat [55] and CPPNs [54] use networks to determine the weight between two neurons as a function of their positions. [35, 34] introduced a similar idea by providing a hierarchical prior for network parameters.

Recurrent Networks and Deep Equilibrium Models. Recurrence and feedback have been shown in psychology and neuroscience to act as modulators or competitive inhibitors to aid feature grouping [21], figure-ground segregation [32] and object recognition [65]. Recurrence-inspired mechanisms also achieve success in feed-forward models. There are two main types of employing recurrence based on if weights are shared across recurrent modules. ResNet [26], a representative of reusing similar structures without weight sharing, introduces parallel residual connections and achieves better performance by going deeper in networks. Similarly, some works [56, 53] also suggest iteratively injecting thus-far representations to the feed-forward network useful. Stacked inference methods [48, 64, 63] are also related while they consider each output in isolation. Some find sharing weights across recurrent modules valuable. They demonstrate applications in temporal modelling [63, 66, 36], spatial attention [44, 9], pose estimation [62, 11], and so on [41, 69]. Such methods usually shine in modeling long-term dependencies. In this work, we recurrently share weights across different layers of a feedback network to reduce network redundancy.

Given stacking weight-shared modules improve the performance, researchers consider running even infinite depth of such modules by making the sequential modules converge to a fixed point [40, 4]. Employing such *equilibrium* models to existing networks, they show improved performance in many natural language processing [4] and computer vision tasks [5, 61]. One issue with deep equilibrium models is that the forward and backward propagation usually takes much more iterations than explicit feed-forward networks. Some work [19] improves the efficiency by making the backward propagation Jacobian free. Another issue is that *infinite* depth and fixed point may not be necessary or even too strict for some tasks. Instead of achieving infinite depth, our model shares

parameters to a certain level. We empirically compare with equilibrium models in Section 5.

Efficient Network Space and Matrix Factorization. Convolution is an efficient and structured matrix-vector multiplication. Arguably, the most fundamental idea in building efficient linear systems is matrix factorization. Given the redundancy in deep convolutional neural network parameters, one can leverage the matrix factorization concept, e.g., factorized convolutions, and design more efficient network classes [30, 33, 57, 51].

3. Recurrent Parameter Generator

Linearly Constrained Neural Optimization. Consider optimizing a network with input data \mathbf{X} , parameters $\hat{\mathbf{W}}$ and loss function L . The optimization can be written as:

$$\min L(\mathbf{X}; \hat{\mathbf{W}}) \text{ s.t. } \hat{\mathbf{W}} = \mathbf{G}\mathbf{W} \text{ (or equally } \mathbf{R}\hat{\mathbf{W}} = 0) \quad (1)$$

where $\hat{\mathbf{W}} = \mathbf{G}\mathbf{W}$ refers to a set of linear constraints, where $\mathbf{G} \in \mathbb{R}^{N \times M}$ is a full-rank tall matrix (i.e. $N \geq M$). Here we refer to $\hat{\mathbf{W}}$ as the constrained parameters and \mathbf{W} as the free parameters. This constraint is a change of variable, i.e., the constrained parameter $\hat{\mathbf{W}}$ is linearly generated from the free parameter \mathbf{W} by generating matrix \mathbf{G} . We can consider \mathbf{W} as a compressed model, which is unpacked into $\hat{\mathbf{W}}$ to construct the large neural network. \mathbf{W} is directly optimized via gradient descent and free to update. In this linearly constrained neural optimization, the model DoF is equivalent to M , which is the dimension of \mathbf{W} . An equivalent form of the constraint $\hat{\mathbf{W}} = \mathbf{G}\mathbf{W}$ is $\mathbf{R}\hat{\mathbf{W}} = 0$, where $\mathbf{R} \in \mathbb{R}^{(N-M) \times N}$ can be derived from SVD of \mathbf{G} .

Recurrent Parameter Generator. Let's assume that we construct a deep convolutional neural network containing L different convolution layers. Let $\mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_L$ be the corresponding L convolutional kernels¹. Rather than using separate sets of parameters for different convolution layers, we create a single set of parameters $\mathbf{W} \in \mathbb{R}^M$ and use it to generate the corresponding parameters $\hat{\mathbf{W}} =$

¹A kernel contains all the filters of one layer. In this paper, we treat each convolutional kernel as a vector. When the kernel is used to do the convolution, it will be reshaped into the corresponding shape.

$[\mathbf{K}_1^T, \mathbf{K}_2^T, \dots, \mathbf{K}_L^T]^T \in \mathbb{R}^N$ for each convolution layer:

$$\mathbf{K}_i = \mathbf{G}_i \cdot \mathbf{W}, i \in \{1, \dots, L\} \quad (2)$$

where \mathbf{G}_i is a fixed predefined generating matrix, which is used to generate \mathbf{K}_i from \mathbf{W} . We call $\mathbf{G} = [\mathbf{G}_1^T, \dots, \mathbf{G}_L^T]^T$ and \mathbf{W} the *recurrent parameter generator* (RPG). In this work, we always assume that the size of \mathbf{W} is not larger than the total parameters of the model, i.e., $|\mathbf{W}| \leq \sum_i |\mathbf{K}_i|$. This means an element of \mathbf{W} will generally be used in more than one layer of a neural network. Additionally, the gradient of \mathbf{W} is a linear superposition of the gradients from each convolution layer. During the neural network training, let's assume convolution kernel \mathbf{K}_i receives gradient $\frac{\partial \ell}{\partial \mathbf{K}_i}$, where ℓ is the loss function. Based on the chain rule, it is clear that the gradient of \mathbf{W} is:

$$\frac{\partial \ell}{\partial \mathbf{W}} = \sum_{i=1}^L \mathbf{G}_i^T \cdot \frac{\partial \ell}{\partial \mathbf{K}_i} \quad (3)$$

Generating Matrices and Destructive Weight Sharing.

There are various ways to create the generating matrices $\{\mathbf{G}_i\}$. While in general \mathbf{G} can be any full-rank tall matrix, this paper focuses on the destructive generating matrices, which are random orthogonal matrices and could prevent different kernels from sharing the representation during weight sharing. Random generating matrices empirically improve the model capacity when the model DoF is fixed. We provide an intuitive theoretical explanation of how random orthogonal matrices prevent representation sharing as follows.

For easier discussion, let us consider a special case, where all of the convolutional kernels have the same size and are used in the same shape in the corresponding convolution layers. The dimension of \mathbf{W} is equal to that of one convolutional layer kernel. In other words, $\{\mathbf{G}_i\}$ are square matrices, and the spatial sizes of all of the convolutional kernels have the same size, $d_{in} \times d_{out} \times w \times h$, and the input channel dimension d_{in} is always equal to the output channel dimension d_{out} . In this case, a filter \mathbf{f} in a kernel can be treated as a vector in \mathbb{R}^{dwh} . Further, we choose \mathbf{G}_i to be a block-diagonal matrix $\mathbf{G}_i = \text{diag}\{\mathbf{A}_i, \mathbf{A}_i, \dots, \mathbf{A}_i\}$, where $\mathbf{A}_i \in O(dwh)$ is an orthogonal matrix that generates each filter of the kernel \mathbf{K}_i from \mathbf{W} , and $O(\cdot)$ denotes the orthogonal group. Similar to the Proposition 2 in [13], we show in the Appendix C that: if $\mathbf{A}_i, \mathbf{A}_j$ are sampled from the $O(dwh)$ Haar distribution and $\mathbf{f}_i, \mathbf{f}_j$ are the corresponding filters (generated by $\mathbf{G}_i, \mathbf{G}_j$ respectively from the same set of entries of \mathbf{W}) from $\mathbf{K}_i, \mathbf{K}_j$ respectively, then we have $E[\langle \mathbf{f}_i, \mathbf{f}_j \rangle] = 0$ and $E\left[\left\langle \frac{\mathbf{f}_i}{\|\mathbf{f}_i\|}, \frac{\mathbf{f}_j}{\|\mathbf{f}_j\|} \right\rangle^2\right] = \frac{1}{dwh}$. Since dwh is usually large, the corresponding filters from $\mathbf{K}_i, \mathbf{K}_j$ are close to orthogonal and generally dissimilar. This shows that even when $\{\mathbf{K}_i\}$ are generated from the same entries of \mathbf{W} , they are prevented from sharing the representation.

Though $\{\mathbf{G}_i\}$ are not updated during training, the size of \mathbf{G}_i can be quite large in general, which can create additional computation and storage overhead. In practice, we can use permutation and element-wise random sign reflection to construct a subset of the orthogonal group as permutations and sign reflections could be implemented with high simplicity and negligible cost. A simple demonstration of $\{\mathbf{G}_i\}$ is demonstrated in Fig.2U². Since pseudo-random numbers are used, it takes only two random seeds to store a random permutation and an element-wise random sign reflection.

Even Parameter Sampling and Model Ring. While it is easy to randomly sample elements from \mathbf{W} when generating parameters for each layer, it may not be optimal as some elements in \mathbf{W} may not be evenly used, and some elements in \mathbf{W} used at all due to sampling fluctuation. A simple equalization technique can be used to guarantee all elements of \mathbf{W} are evenly sampled. Suppose the size of \mathbf{W} is M , and the size of parameter $\hat{\mathbf{W}}$ of the model to be generated is N , $N > M$. As we mentioned earlier, there are L layers and they require $\{\|K_1\|, \dots, \|K_L\|\}$ parameters respectively. As $N > M$, we can use $\hat{\mathbf{W}}$ as a ring: we first draw the first $\|K_1\|$ parameters from $\hat{\mathbf{W}}$ followed by a pre-generated random permutation p_1 and a pre-generated random element-wise sign flipping b_1 to construct layer-1 kernel \mathbf{K}_1 . Then we draw the next $\|K_2\|$ parameters from $\hat{\mathbf{W}}$ followed by pre-generated random permutation p_2 and a pre-generated random element-wise sign flipping b_2 . We continue this process and wrap around when there is not enough entries left from $\hat{\mathbf{W}}$. We refer to $\hat{\mathbf{W}}$ together with this sampling strategy as *model rings* since the free parameters are recurrently used in a loop. We illustrate the general parameter generator in Fig.2U and RPG in Fig.2L. This For data saving efficiency, we just need to save several random seed numbers instead of saving the pre-generated permutations $\{p_1, \dots, p_L\}$ and sign flipping operations $\{b_1, \dots, b_L\}$.

Batch Normalization. Model performance is relatively sensitive to the batch normalization parameters. For better performance, each convolution layer needs to have its own batch normalization parameters. In general, however, the size of batch normalization is relatively negligible. Yet when \mathbf{W} is extremely small (e.g., 36K parameters), the size of batch normalization should be considered.

4. RPG at Multiple Scales

We discuss the general idea of parameter generators where only one RPG is shared globally across all layers previously. We could also create several local RPGs, each of which is shared at certain scales, such as blocks and sub-networks. Such RPGs may be useful for certain applications such as recurrent modeling.

²Permutations and element-wise random sign reflection conceptually are subgroups from the orthogonal group, but we shall never use them in the matrix form for the obvious efficiency purpose.

RPGs at Block-Level. Many existing network architectures reuse the same design of network blocks multiple times for higher learning capacity, as discussed in the related work. Instead of using one global RPG for the entire network, we could alternatively create several RPGs that are shared within certain network blocks. We take Res18 [26] as a concrete example. Res18 has four building blocks. Every block has 2 residual convolution modules. We create four local RPGs for Res18. Each RPG is shared within the corresponding building block, where the size of the RPG is flexible and can be determined by users. Fig.3M) illustrates how RPGs can be shared at the block-level.

RPGs at Sub-Network-Level. Reusing sub-networks, or recurrent networks, has achieved success in many tasks as they iteratively refine and improve the prediction. Parameters are often shared when reusing the sub-networks. This may not be optimal as sub-networks at different stages iteratively improve the prediction, and shared parameters may limit the learning capacity at different stages. However, not sharing parameters at all greatly increases the model size. RPG can be created for each sub-network. Such design leads to a much smaller DoF, while parameters of different sub-networks are orthogonal by undergoing destructive changes. We show applications of sub-network-level RPGs for pose estimation and multitask regression (Section 5.3 and 5.4). Fig.3R) illustrates sub-network-level RPGs.

5. Experimental Results

We evaluate the performance of RPG with various tasks illustrated in Fig.3. For classification, RPG was used for the entire network except for the last fully-connected layer. We discuss performance with regard to *backbone DoF*, the actual number of parameters of the backbone. For example, Res18 has 11M backbone parameters and 512K fc parameters, and RPG was applied to reduce 11M backbone DoF only.

5.1. CIFAR Classification

Implementation Details. CIFAR experiments use 128 batch size, $5e-4$ weight decay, initial learning rate of 0.1 with gamma of 0.1 at epoch 60, 120 and 160. We use Kaiming initialization [25] with adaptive scaling. Shared parameters are initialized with a particular variance and scale the parameters for each layer to make it match the Kaiming initialization.

Compared to Deep Equilibrium Models. As a representative of implicit models, deep equilibrium models [4] reduce model DoF by finding fix points via additional optimizations. We compare the image classification accuracy on CIFAR10 and CIFAR100, as well as the inference time on CIFAR100 (Table 1). Following the settings of MDEQ [5], an image was sequentially fed into the initial convolutional block, the multi-scale deep equilibrium block (dubbed as *MS* block), and the classification head. MDEQ [5] achieves *infinite* MS blocks by finding the fixed point of the MS block. We reuse




Figure 4: **a)** Large models are known to have high redundancy and low degree of freedom (DoF). They could be pruned to small models, e.g. high filter similarity of different layers in VGG16 is observed. **b)** Ablation studies of permutation and sign reflection of Res34-RPG. Having both matrices gives the highest performance.

the MS block two to four times without increasing the model DoF. RPG achieves 3% - 6% gain on CIFAR10 and 3% - 6% gain on CIFAR100. RPG inference time is 15 - 25 times smaller than MDEQ since MDEQ needs additional time to solve equilibrium during training.

Global RPG with Varying Model DoF. We create one global RPG to generate parameters for convolution layers of ResNet and refer to it as *ResNet-RPG*. We report CIFAR100 top-1 accuracy of ResNet-RPG18 and ResNet-RPG34 at different model DoF (Table 3 and Fig.6 in Appendix B). Compared to ResNet, ResNet-RPG achieves higher accuracy at the same model DoF. Specifically, we achieve 36% CIFAR100 accuracy with only 8K backbone DoF. Further, ResNet34-RPG achieves higher accuracy than ResNet18-RPG, indicating increasing time complexity gives performance gain. We observe log-linear DoF-accuracy relationship, with details in *Power Law* of the following subsection.

Local RPGs at the Block-Level. In the previous Res-RPG

Table 1: RPG compared with multiscale deep equilibrium models (MDEQ) [5] on CIFAR10 and CIFAR100 classification. At the same number of model DoF, RPG achieves 3% - 6% performance gain with 15 - 25x less inference time. Inference time is measured by milliseconds per image.

Accuracy (%)	MDEQ	Our RPG (same DoF)		
		2x MS blk	3x MS blk	4x MS blk
CIFAR10	85.1	88.5	90.1	90.9
CIFAR100	59.8	62.8	64.7	65.7
Inference time (ms)	3.15	0.12	0.18	0.22

Table 2: ResNet-RPG outperforms existing DoF reduction methods [23, 12, 67] on CIFAR100. Additionally, a global RPG outperforms block-wise local RPGs.

	DoF	Acc. (%)
R18-vanilla	11M	77.5
R34-RPG.blk	11M	78.5
R34-RPG	11M	78.9
R34-random weight share	11M	74.9
R34-DeepCompression [23]	11M	72.2
R34-Hash [12]	11M	75.6
R34-Lego [67]	11M	78.4
R34-vanilla	21M	79.1

Table 3: ResNet-RPG consistently achieves higher performance at the same model DoF. We report ImageNet and CIFAR100 top-1 accuracy and backbone DoF for ResNet-vanilla and ResNet-RPG.

Acc. (%)	R18-RPG			R18-vanilla	R34-RPG			R34-vanilla
ImageNet	40.0	67.2	70.5	70.5	41.6	69.1	73.4	73.4
CIFAR100	60.2	75.6	77.6	77.6	61.7	76.5	78.9	79.1
Model DoF	45K	2M	5.5M	11M	45K	2M	11M	21M

experiments, we use one global RPG for the entire network. We also evaluate the performance when RPGs are shared locally at a block level, as discussed in Section 5.4. In Table 2, compared to plain ResNet18 at the same DoF, our block-level RPG network gives 1.0% gain. In contrast, our ResNet-RPG (parameters are evenly distributed) gives a 1.4% gain. Using one global RPG where parameters of each layer are evenly distributed is 0.4% higher than multiple RPGs.

Comparison to Baselines. Table 2 compares RPG and other model DoF reduction methods including random weight sharing, weight sharing with the deep compression [23], hashing trick [12] and weight sharing with Lego filters [67]. We also compare with HyperNetworks [22] in Appendix D. At the same model DoF, RPG outperforms all other baselines, demonstrating the effectiveness of the proposed method.

RPG for Transformers. We apply RPG for a vision transformer ViT [17] and report results in Fig.5a. Specifically, the ViT-tiny model with 6 transformer layers, 4 attention heads and 64 embedding dimensions, is used as a baseline. A log-linear relationship is also identified in ViT-RPG.

5.2. ImageNet Classification

Implementation Details. All ImageNet experiments use batchsize of 256, weight decay of $3e-5$, and an initial learning rate of 0.3 with gamma of 0.1 every 75 epochs and 225 epochs in total. Our schedule is different from the standard schedule as the weight-sharing mechanism requires different training dynamics. We tried a few settings and found this one to be the best for RPG.

RPG with Varying Model DoF. We use RPG with different DoF for ResNet and report the top-1 accuracy (Table 3 and Fig.1e)). ResNet-RPGs consistently achieve higher performance than ResNets under the same model DoF. Specifically, ResNet-RPG34 achieves the same accuracy 73.4% as ResNet34 with only half of ResNet34 backbone DoF. ResNet-RPG18 also achieves the same accuracy as ResNet18 with only half of ResNet18 backbone DoF. Further, RPG networks have higher generalizability (Section 5.6).

Power Law. Empirically, accuracy and model DoF follow a power law, when RPG DoF is lower than 50% ResNet-vanilla DoF (Fig.1d). The exponents of the power laws are the same for ResNet18-RPG and ResNet34-RPG on ImageNet. The scaling law may be useful for estimating the network accuracy without training the network. Similarly, [29] also identifies a power law for accuracy and model DoF of transformers. The proposed RPG enables *under-parameterized* models for large-scale datasets such as ImageNet, which may unleash more new studies and findings.

5.3. Pose Estimation

Implementation Details. We superpose sub-networks for pose estimation with a globally shared RPG. Hourglass networks [46] are used as the backbone. An input image is first fed to an initial convolution block to obtain a feature map, which is then fed to multiple stacked pose estimation sub-networks. Each sub-network outputs a pose estimation prediction, which is penalized by the pose estimation loss. Convolutional pose machine (CPM) [62] share all sub-networks weights. We create one global RPG to generate parameters for each sub-network. Our model size is set to the same as CPM. We also compare with larger models where parameters of sub-networks are not shared.

We evaluate on MPII Human Pose dataset [2], a benchmark for articulated human pose estimation, which consists of over 28K training samples over 40K people with annotated body joints. We use the hourglass network [46] as backbone and follow all their settings.

Results and Analysis. We report the Percentage of Correct Key-points at 50% threshold (PCK@0.5) of different methods in Table 4. CPM [62] share all parameters for different sub-networks. We use one RPG that is shared globally at the same size as CPM. For reference, we also compare with the no-sharing model as the performance ceiling. Adding the number of recurrences leads to performance gain for all methods. At the same model size, RPG achieves higher PCK@0.5 compared to CPM. Increasing the number of parameters by not sharing sub-network parameters also leads to some performance gain.

5.4. Multi-Task Regression

Implementation Details. We superpose sub-networks for multi-task regression with multiple RPGs at the building-block level. We focus on predicting depth and normal maps

Table 4: RPG outperforms CPM [62] at the same DoF. We report pose estimation performance (model DoF) on MPII human pose compared with CPM [62]. The metric is PCKh@0.5.

Acc. (DoF)	CPM [62]	RPG	No shared w.
1x sub-net	84.7 (3.3M)		
2x sub-nets	86.1 (3.3M)	86.5 (3.3M)	87.1 (6.7M)
4x sub-nets	86.5 (3.3M)	87.3 (3.3M)	88.0 (13.3M)

Table 5: RPG achieves the best accuracy without sharing batch normalize parameters and with permutation and sign reflection. We report multitask regression errors on S3DIS with sub-net architecture as [49]. Lower is better. All methods share the same DoF. Sub-net is reused once.

RMSE (%)	Depth	Normal
Vanilla model	25.5	41.0
RPG with shared BN	24.7	40.3
Reuse & new BN	24.0	39.4
Reuse & new BN & perm. and reflect.	22.8	39.1




Figure 5: **a)** A log-linear DoF-accuracy relationship exists for RPGs applied to vision transformer ViT [17]. **b)** RPG converges faster than the vanilla model. We plot the CIFAR10 accuracy (smoothed by moving average) versus training iterations for Res18-vanilla and Res18-RPG. RPG converges at 1k iterations while the vanilla model converges at 1.7k. **c)** RPG consistently converges faster. The reduction becomes substantial with the increasing batchsize, e.g., at batchsize 1024, RPG takes 41% less iterations to converge. Denote final accuracy as P_f , the convergence iteration is defined when current smoothed accuracy (by moving average) is within 5% range of P_f .

Table 6: RPG achieves higher post-pruning CIFAR10 accuracy and similar post-pruning accuracy drops as SOTA fine-grained pruning approach IMP [18]. Fine-grained pruning is used for reducing DoF.

	acc before	acc after ↓ DoF	acc drop	model DoF
R18-IMP [18]	92.3	90.5	1.8	274k
R18-RPG	95.0	93.0	2.0	274k

Table 7: RPG achieves similar post-pruning ImageNet performance as SOTA coarse-grained approach Knapsack [1] at the same FLOPs. Coarse-grained pruning is used for reducing RPG FLOPs.

	DoF before pruning	Pruned acc.	FLOPs
R18-Knapsack	11.2M	69.35%	1.09e9
Pruned R18-RPG	5.6M	69.10%	1.09e9

from a given image. We stack multiple SharpNet [49], a network for monocular depth and normal estimation. Specifically, we create multiple RPGs at the SharpNet building-block level. That is, parameters of corresponding blocks of different sub-networks are generated from the same RPG.

We evaluate the monocular depth and normal prediction performance on a 3D indoor scene dataset [3], which contains over 70K images with corresponding depths and normals covering over 6,000 m² indoor area. We follow all settings of SharpNet [49], a SOTA monocular depth and normal estimation method.

Results and Analysis. We report the mean square errors for depth and normal estimation in Table 5. Compared to one-time inference without recurrence, our RPG network gives 3% and 2% gain for depth and normal estimation, respectively. Directly sharing weights but using new batch normalization layers decrease the performance by 1.2% and 0.3% for depth and normal. Sharing weights and normalization layers further decrease the performance by 0.7% and 0.9% for depth and normal.

5.5. Pruning RPG

Fine-Grained Pruning. Fine-grained pruning methods aim to reduce the model DoF by sparsifying weight matrices.

Such methods usually do not reduce the inference speed, although custom algorithms [20] may improve the speed. At the same model DoF, RPG outperforms state-of-the-art fine-grained pruning method IMP [18]. Accuracy drops of RPG and IMP are similar, both around 2% (Table 6). It is worth noting that although IMP has no run time improvement in regular settings, it could save inference time with customized sparse GPU kernels [20].

Coarse-Grained Pruning. While RPG is not designed to reduce FLOPs, it can be combined with coarse-grained pruning to reduce FLOPs. We prune RPG filters with the lowest ℓ_1 norms. Table 7 shows that the pruned RPG achieves on-par performance as state-of-the-art coarse-grained pruning method Knapsack [1] at the same FLOPs.

5.6. Analysis

Convergence rate. Compared with the vanilla model, RPG optimizes in a parameter subspace $\tilde{\mathbf{W}} = \mathbf{G}\mathbf{W}$ with fewer DoF. Would such constrained optimization lead to a faster convergence rate? We analyze the convergence rate of Res18-vanilla and Res18-RPG (DoF is 5.5M, 50% of the vanilla model) with different batchsizes. All models are trained with multi-step SGD optimizer and they all reach > 94.1% final CIFAR10 accuracy. For simplicity, we analyze the first optimization stage where learning rate has not decayed.

Fig.5b plots the accuracy (smoothed with moving averages) v.s. training iterations with batchsize 1024. RPG has a faster convergence rate than vanilla models. We also analyze the smoothed accuracy and identify the convergence iteration versus batchsize in Fig.5c. RPG consistently converges faster than the vanilla model, and the reduction becomes substantial with the increasing batchsize.

Comparison to Model Compression Methods. We report ResNet-RPG performance with different model DoF and existing compression methods on ImageNet (Fig.1e). RPG networks outperform SOTA methods such as [1, 16, 28, 27, 15, 37]. For example, at the same model DoF, our RPG

Table 8: RPG increases the model generalizability. **(a)** ResNet-RPG has lower training-validation accuracy gap on ImageNet classification. The metric is training accuracy minus validation accuracy. Lower is better. **(b)** Using RPG for pose estimation also decreases the training and validation performance GAP. The metric is training PCK@0.5 minus validation PCK@0.5. Lower is better. **(c)** ResNet with RPG has higher performance on out-of-distribution dataset ObjectNet [6]. The model is trained on ImageNet only and directly evaluated on ObjectNet.

(a) IN train-val gap			(b) Pose train-val gap			(c) OOD on ObjectNet				
Acc gap (%)	vanilla	RPG	Acc gap (%)	no shared w	shared w	RPG		R18	R34-RPG	R34
R18	-0.7	-2.7	2x sub-nets	1.15	1.13	0.64	DoF	11M	11M	21M
R34	1.1	-2.3	4x sub-nets	1.98	1.70	1.15	Acc. (%)	13.4	16.5	16.0

network has 0.6% gain over the knapsack pruning [1], a SOTA method of ImageNet pruning.

Storage. RPG models only need to save the effective parameter \mathbf{W} , which has the size of the model DoF, since the generation matrix G is saved as a random seed at no cost. The storage space of the model file can be diminished to satisfy a smaller storage limit for inference and a faster model file transfer. Empirically on PyTorch platform, ResNet18-vanilla model file is 45MB. With no accuracy loss, ResNet18-RPG model save file size is 23MB (\downarrow 49%). With 2 percentage point accuracy loss, RPG save file size is 9.5MB (\downarrow 79%).

Generalizability. We report the performance gap between training and validation set on ImageNet (Table 8(a)) and MPII pose estimation (Table 8(b)). CPM [62] serves as the baseline pose estimation method. RPG models consistently achieve lower gaps between training and validation sets, indicating the RPG model suffers less from over-fitting.

We also report the out-of-distribution performance of RPG models. ObjectNet [6] contains 50k images with 113 classes overlapping with ImageNet. Existing models are reported to have a large performance drop on ObjectNet. We directly evaluate the performance of ImageNet-trained model on ObjectNet without any fine-tuning (Table 8(c)). With the same backbone DoF, R18-RPG achieves a 3% gain compared to R18-vanilla. With the same network architecture design, R34-RPG achieves 0.5% gain compared to R34. This indicates RPG networks have higher out-of-distribution performance even with smaller model DoF.

Quantization. Network quantization can reduce model size with minimal accuracy drop. It is of interest to study if RPG models, whose parameters have been shrunk, can be quantized. After 8-bit quantization, the accuracy of ResNet18-RPG (5.6M DoF) only drop 0.1 percentage point on ImageNet, indicating RPG can be quantized for further model size reduction. Details are in Appendix A.

Security. Permutation matrices generated by the random seed can be considered as security keys to decode the model. Further, only random seeds to generate generating matrix G need to be saved and transferred at negligible cost.

5.7. Ablation Studies

We conduct ablation studies on CIFAR100 to analyze functions of permutation and reflection matrices (Fig.4b). We evaluate ResNet-RPG34 with 2M backbone DoF. Permutation and sign reflection together achieves 76.5% accuracy,

while permutation only achieves 75.8%, and sign reflection only achieves 71.1%. Training with neither permutation nor reflection matrices achieves 70.7%. This suggests permutation and sign reflection matrices increase RPG performance.

6. Discussion

The common practice in neural network compression is to prune weights from a trained large model with many parameters or degrees of freedom (DoF). Our key insight is that a direct and drastically different approach might work faster and better: We start from a lean model with a small DoF, which can be linearly unpacked into a large model with many parameters. Then we can let the gradient descent automatically find the best model under the linear constraints. Our work is a departure from mainstream approaches towards model optimization and parameter reduction. We show how the model DoF and actual parameter size can be decoupled: we can define an arbitrary network of an arbitrary DoF.

We limit our scope to optimization with random linear constraints, termed destructive weight sharing. However, in general, there might also exist nonlinear RPGs and efficient nonlinear generation functions to create convolutional kernels from a shared model ring \mathbf{W} . Further, although RPG focuses on reducing model DoF, it can be quantized and pruned to further reduce the FLOPs and runtime.

To sum up, we develop an efficient approach to build an arbitrarily complex neural network with any amount of DoF via a recurrent parameter generator. On a wide range of applications, including classification, pose estimation and multitask regression, we show RPG consistently achieves higher performance at the same model DoF. Further, we show such networks converge faster, are less likely to overfit and have higher performance on out-of-distribution data.

RPG can be added to any existing network flexibly with any amount of DoF at the user’s discretion. It provides new perspectives for recurrent models, equilibrium models, and model compression. It also serves as a tool for understanding relationships between network properties and network DoF by factoring out the network architecture.

References

- [1] Yonathan Aflalo, Asaf Noy, Ming Lin, Itamar Friedman, and Lihi Zelnik. Knapsack pruning with inner distillation. *arXiv preprint arXiv:2002.08258*, 2020.

- [2] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3686–3693, 2014.
- [3] Iro Armeni, Sasha Sax, Amir R Zamir, and Silvio Savarese. Joint 2d-3d-semantic data for indoor scene understanding. *arXiv preprint arXiv:1702.01105*, 2017.
- [4] Shaojie Bai, J Zico Kolter, and Vladlen Koltun. Deep equilibrium models. *Advances in Neural Information Processing Systems*, 32:690–701, 2019.
- [5] Shaojie Bai, Vladlen Koltun, and J Zico Kolter. Multiscale deep equilibrium models. *Advances in Neural Information Processing Systems*, 33, 2020.
- [6] Andrei Barbu, David Mayo, Julian Alverio, William Luo, Christopher Wang, Dan Gutfreund, Josh Tenenbaum, and Boris Katz. Objectnet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. *Advances in neural information processing systems*, 32:9453–9463, 2019.
- [7] Davis Blalock, Jose Javier Gonzalez Ortiz, Jonathan Frankle, and John Guttag. What is the state of neural network pruning? In *Proceedings of Machine Learning and Systems*, 2020.
- [8] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020.
- [9] Nicholas J Butko and Javier R Movellan. Optimal scanning for faster object detection. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2751–2758. IEEE, 2009.
- [10] Han Cai, Ligeng Zhu, and Song Han. Proxylessnas: Direct neural architecture search on target task and hardware. In *International Conference on Learning Representations*, 2018.
- [11] Joao Carreira, Pulkit Agrawal, Katerina Fragkiadaki, and Jitendra Malik. Human pose estimation with iterative error feedback. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4733–4742, 2016.
- [12] Wenlin Chen, James Wilson, Stephen Tyree, Kilian Weinberger, and Yixin Chen. Compressing neural networks with the hashing trick. In *International conference on machine learning*, pages 2285–2294. PMLR, 2015.
- [13] Brian Cheung, Alex Terekhov, Yubei Chen, Pulkit Agrawal, and Bruno Olshausen. Superposition of many models into one. In *Advances in neural information processing systems*, 2019.
- [14] Piotr Dollár, Mannat Singh, and Ross Girshick. Fast and accurate model scaling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 924–932, 2021.
- [15] Xuanyi Dong, Junshi Huang, Yi Yang, and Shuicheng Yan. More is less: A more complicated network with less inference complexity. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5840–5848, 2017.
- [16] Xuanyi Dong and Yi Yang. Network pruning via transformable architecture search. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [17] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [18] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M Roy, and Michael Carbin. Stabilizing the lottery ticket hypothesis. *arXiv preprint arXiv:1903.01611*, 2019.
- [19] Samy Wu Fung, Howard Heaton, Qiuwei Li, Daniel McKenzie, Stanley Osher, and Wotao Yin. Fixed point networks: Implicit depth models with jacobian-free backprop. *arXiv preprint arXiv:2103.12803*, 2021.
- [20] Trevor Gale, Matei Zaharia, Cliff Young, and Erich Elsen. Sparse GPU kernels for deep learning. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2020*, 2020.
- [21] Charles D Gilbert and Mariano Sigman. Brain states: top-down influences in sensory processing. *Neuron*, 54(5):677–696, 2007.
- [22] David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016.
- [23] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. In *Proceedings of the International Conference on Learning Representations*, 2016.
- [24] Yongchang Hao, Shilin He, Wenxiang Jiao, Zhaopeng Tu, Michael Lyu, and Xing Wang. Multi-task learning with shared encoder for non-autoregressive machine translation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3989–3996, 2021.
- [25] 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.
- [26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [27] Yang He, Guoliang Kang, Xuanyi Dong, Yanwei Fu, and Yi Yang. Soft filter pruning for accelerating deep convolutional neural networks. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 2234–2240, 2018.
- [28] Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4340–4349, 2019.

- [29] Tom Henighan, Jared Kaplan, Mor Katz, Mark Chen, Christopher Hesse, Jacob Jackson, Heewoo Jun, Tom B Brown, Prfulla Dhariwal, Scott Gray, et al. Scaling laws for autoregressive generative modeling. *arXiv preprint arXiv:2010.14701*, 2020.
- [30] 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.
- [31] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Quantized neural networks: Training neural networks with low precision weights and activations. *The Journal of Machine Learning Research*, 18(1):6869–6898, 2017.
- [32] JM Hupé, AC James, BR Payne, SG Lomber, P Girard, and J Bullier. Cortical feedback improves discrimination between figure and background by v1, v2 and v3 neurons. *Nature*, 394(6695):784–787, 1998.
- [33] 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.
- [34] Theofanis Karaletsos and Thang D. Bui. Hierarchical gaussian process priors for bayesian neural network weights. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [35] Theofanis Karaletsos, Peter Dayan, and Zoubin Ghahramani. Probabilistic meta-representations of neural networks. *arXiv preprint arXiv:1810.00555*, 2018.
- [36] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3128–3137, 2015.
- [37] Ashish Kheta and Zohar Karnin. Prunenet: Channel pruning via global importance. *arXiv preprint arXiv:2005.11282*, 2020.
- [38] Anders Krogh and John A Hertz. A simple weight decay can improve generalization. In *Advances in neural information processing systems*, pages 950–957, 1992.
- [39] Yann LeCun, John S Denker, and Sara A Solla. Optimal brain damage. In *Advances in neural information processing systems*, pages 598–605, 1990.
- [40] Yann LeCun, D Touresky, G Hinton, and T Sejnowski. A theoretical framework for back-propagation. In *Proceedings of the 1988 connectionist models summer school*, volume 1, pages 21–28, 1988.
- [41] Ke Li, Bharath Hariharan, and Jitendra Malik. Iterative instance segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3659–3667, 2016.
- [42] Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning. In *International Conference on Learning Representations*, 2018.
- [43] C Louizos, M Reisser, T Blankevoort, E Gavves, and M Welling. Relaxed quantization for discretized neural networks. In *International Conference on Learning Representations*, ICLR, 2019.
- [44] Volodymyr Mnih, Nicolas Heess, Alex Graves, and Koray Kavukcuoglu. Recurrent models of visual attention. In *Advances in Neural Information Processing Systems*, 2014.
- [45] Michael C Mozer and Paul Smolensky. Using relevance to reduce network size automatically. *Connection Science*, 1(1):3–16, 1989.
- [46] 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.
- [47] Steven J. Nowlan and Geoffrey E. Hinton. Simplifying neural networks by soft weight-sharing. *Neural Computation*, 4(4):473–493, 1992.
- [48] Varun Ramakrishna, Daniel Munoz, Martial Hebert, James Andrew Bagnell, and Yaser Sheikh. Pose machines: Articulated pose estimation via inference machines. In *European Conference on Computer Vision*, pages 33–47. Springer, 2014.
- [49] Michael Ramamonjisoa and Vincent Lepetit. Sharpnet: Fast and accurate recovery of occluding contours in monocular depth estimation. *The IEEE International Conference on Computer Vision (ICCV) Workshops*, 2019.
- [50] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European conference on computer vision*, pages 525–542. Springer, 2016.
- [51] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- [52] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- [53] Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. Highway networks. *arXiv preprint arXiv:1505.00387*, 2015.
- [54] Kenneth O Stanley. Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines*, 8(2):131–162, 2007.
- [55] Kenneth O Stanley, David B D’Ambrosio, and Jason Gauci. A hypercube-based encoding for evolving large-scale neural networks. *Artificial life*, 15(2):185–212, 2009.
- [56] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [57] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114, 2019.
- [58] Alvin Wan, Xiaoliang Dai, Peizhao Zhang, Zijian He, Yuan-dong Tian, Saining Xie, Bichen Wu, Matthew Yu, Tao Xu, Kan Chen, et al. Fbnetv2: Differentiable neural architecture

- search for spatial and channel dimensions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12965–12974, 2020.
- [59] Li Wan, Matthew Zeiler, Sixin Zhang, Yann Le Cun, and Rob Fergus. Regularization of neural networks using dropconnect. In *International conference on machine learning*, pages 1058–1066. PMLR, 2013.
- [60] Jiayun Wang, Yubei Chen, Rudrasis Chakraborty, and Stella X Yu. Orthogonal convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11505–11515, 2020.
- [61] Tiancai Wang, Xiangyu Zhang, and Jian Sun. Implicit feature pyramid network for object detection. *arXiv preprint arXiv:2012.13563*, 2020.
- [62] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional pose machines. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 4724–4732, 2016.
- [63] David Weiss and Benjamin Taskar. Structured prediction cascades. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 916–923. JMLR Workshop and Conference Proceedings, 2010.
- [64] David H Wolpert. Stacked generalization. *Neural networks*, 5(2):241–259, 1992.
- [65] Dean Wyatte, Tim Curran, and Randall O’Reilly. The limits of feedforward vision: Recurrent processing promotes robust object recognition when objects are degraded. *Journal of Cognitive Neuroscience*, 24(11):2248–2261, 2012.
- [66] Shi Xingjian, Zhouong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*, pages 802–810, 2015.
- [67] Zhaohui Yang, Yunhe Wang, Chuanjian Liu, Hanting Chen, Chunjing Xu, Boxin Shi, Chao Xu, and Chang Xu. Legonet: Efficient convolutional neural networks with lego filters. In *International Conference on Machine Learning*, pages 7005–7014. PMLR, 2019.
- [68] Jiahui Yu, Linjie Yang, Ning Xu, Jianchao Yang, and Thomas Huang. Slimmable neural networks. In *International Conference on Learning Representations*, 2018.
- [69] Amir R Zamir, Te-Lin Wu, Lin Sun, William B Shen, Bertram E Shi, Jitendra Malik, and Silvio Savarese. Feed-back networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1308–1317, 2017.
- [70] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. In *ICLR*, 2017.

Appendices

We first show RPG networks could be quantized with minimal accuracy drop for compression purpose in Section A. We then provide a figure revealing log-linear DoF-accuracy relationship in Section B. We also provide proof for the orthogonal proposition in the main paper (Section C). Finally, we provide detailed comparison and discussion to a closely related work HyperNetworks [22] in Section D.

Additionally, we provide the most important code to reproduce the layer superposition experiments on ImageNet in supplementary as a .tgz file. The rest of code is also ready for release, and will be released after additional internal review.

A. Quantize RPG

Quantization refers to techniques for performing computations and storing tensors at lower bitwidths than floating point precision. Quantization can reduce model size with tiny accuracy drop. Table 9 shows that with 8-bit quantization, ResNet18-vanilla has an accuracy drop of 0.3 percentage point, while our ResNet18-RPG has an accuracy drop of 0.1 percentage point. RPG models can be quantized for further model size reduction with a negligible accuracy drop.

Table 9: RPG model can be quantized with very tiny accuracy drop. With 8-bit quantization on ImageNet, ResNet18-vanilla has an accuracy drop of 0.3 percentage point, while our ResNet18-RPG has an accuracy drop of 0.1 percentage point.

	# Params	Acc before	Acc after ↓ quantization	Acc drop
R18-vanilla	11M	69.8	69.5	0.3
R18-RPG	5.6M	70.2	70.1	0.1

B. CIFAR100 Accuracy versus DoF

Fig.6 plots CIFAR100 classification accuracy versus model DoF. We observe a similar log-linear relationship as in ImageNet.

C. Proof to the Orthogonal Proposition

We provide proofs to the orthogonal proposition mentioned in Section 3 of the main paper. Suppose we have two vectors $\mathbf{f}_i = \mathbf{A}_i \mathbf{f}$, $\mathbf{f}_j = \mathbf{A}_j \mathbf{f}$, where $\mathbf{A}_i, \mathbf{A}_j$ are sampled from the $O(M)$ Haar distribution.

Proposition 1. $E[\langle \mathbf{f}_i, \mathbf{f}_j \rangle] = 0$.




Figure 6: Log-linear DoF-accuracy relationship of CIFAR100 accuracy and model DoF on CIFAR100. RPG achieves the same accuracy as vanilla ResNet with 50% DoF.

Proof.

$$\begin{aligned}
 E[\langle \mathbf{f}_i, \mathbf{f}_j \rangle] &= E[\langle \mathbf{A}_i \mathbf{f}, \mathbf{A}_j \mathbf{f} \rangle] \\
 &= E[\langle \mathbf{A}_i \mathbf{f}, \mathbf{A}_j \mathbf{f} \rangle] \\
 &= E[\langle \mathbf{f}, \mathbf{A}_i^T \mathbf{A}_j \mathbf{f} \rangle] \\
 &= \mathbf{f}^T E[\mathbf{A}_i^T \mathbf{A}_j] \mathbf{f} \\
 &= 0
 \end{aligned}$$

where $\mathbf{A}_i^T \mathbf{A}_j$ is equivalently a random sample from $O(M)$ Haar distribution and its expectation is clearly 0. \square

Proposition 2. $E[\langle \frac{\mathbf{f}_i}{\|\mathbf{f}_i\|}, \frac{\mathbf{f}_j}{\|\mathbf{f}_j\|} \rangle^2] = \frac{1}{M}$.

Proof.

$$\begin{aligned}
 E\left[\left\langle \frac{\mathbf{f}_i}{\|\mathbf{f}_i\|}, \frac{\mathbf{f}_j}{\|\mathbf{f}_j\|} \right\rangle^2\right] &= \frac{E[\langle \mathbf{A}_i \mathbf{f}, \mathbf{A}_j \mathbf{f} \rangle^2]}{\|\mathbf{f}\|_2^2 \|\mathbf{f}\|_2^2} \\
 &= E\left[\left\langle \mathbf{A} \frac{\mathbf{f}}{\|\mathbf{f}\|}, \frac{\mathbf{f}}{\|\mathbf{f}\|} \right\rangle^2\right],
 \end{aligned}$$

where $\mathbf{A} = \mathbf{A}_i^T \mathbf{A}_j \sim O(M)$ Haar distribution

Due to the symmetry,

$$= E\left[\left\langle \mathbf{A} \frac{\mathbf{f}}{\|\mathbf{f}\|}, (1, 0, 0, \dots, 0)^T \right\rangle^2\right]$$

Let $\mathbf{g} = \mathbf{A} \frac{\mathbf{f}}{\|\mathbf{f}\|}$,

$$\begin{aligned}
 &= E[g_1^2] \\
 &= \frac{1}{M}
 \end{aligned}$$

since \mathbf{g} is a random unit vector and $E[\sum_{k=1}^M g_k^2] = \sum_{k=1}^M E[g_k^2] = 1$. \square

D. Comparison to HyperNetworks

HyperNetworks [22] share similarity with RPG as both methods reduce model DoF. Specifically, HyperNetworks rely on learnable modules to generate network parameters. We compare with them and report results in Table 10. On CIFAR100 with the embedding dimension of 64 and the same model size, HyperNetworks has 68x FLOPs as our RPG, yet 10 percentage points lower than RPG in accuracy.

Table 10: RPG outperforms HyperNetworks [22] with same DoF on CIFAR100. HyperNetworks has 68x FLOPs as our RPG, yet 10 percentage points lower than RPG in accuracy.

	model DoF	FLOPs	CIFAR100 Acc.
HyperNet [22]	632k	2.49G	61.3%
RPG	632k	36.7M	71.6%

RPG can be considered as an extreme and minimal version of HyperNetworks, one without a network. However, RPG’s unique design and implementation delivers the following advantages over HyperNetworks:

1. HyperNetworks add substantial FLOPs to the network and render it less practical. Given a network architecture, RPG adds minimal to no additional computation, as the permutation and sign reflection can be efficiently implemented. However, HyperNetworks use a weight generation network to generate the primary network weights. A hypernet mainly uses matrix multiplication and introduces substantial FLOPs. In the table below, we analyze FLOPs of HyperNetwork for ResNet18 with the embedding dimension of 64. FLOPs of a vanilla-Res18 for ImageNet (224 input size) and CIFAR100 (32 input size) are 1.8G and 36.7M, whereas the weight generation part of the HyperNet-Res18 takes 2.45G FLOPs. This means the weight generation FLOPs are 1.4 times of vanilla-Res18 for ImageNet and 67 times of that of CIFAR100. Empirically, we find the training and inference time HyperNet-Res18 is around 70x larger than vanilla-Res18.
2. HyperNetworks do not have an arbitrary DoF (number of reduced parameters). RPG uses a model ring of a size (model DoF) that can be arbitrarily determined. In HyperNetworks, the weight generation network uses the same hyper-weight and requires embedding to be of a certain size so that the matrix multiplication can be used for generating primary network weights. Therefore, the model DoF or reduced number of parameters cannot be arbitrarily determined. In other words, RPG decouples the model DoF (actual parameters) and the network architecture, while HyperNetworks have model DoF and

architecture tightly coupled together, a highly restrictive limitation.

3. Weights generated by HyperNetworks may be coupled and not optimized for different layers. HyperNetworks use only one weight generation network parameterized by hyper-weight to generate all primary network weights. This may not be optimal as different layers of the primary network may need different weight generation networks. Additionally, matrix multiplication is used for generating weights, and the generated primary network weights may be coupled. On the other hand, RPG has destructive weight sharing, which improves the network performance by decoupling cross-layer network weights. We will add these results and discussions in the revision to clarify the differences between RPG and HyperNetworks.