

# RANP: Resource Aware Neuron Pruning at Initialization for 3D CNNs

Zhiwei Xu<sup>1,3</sup> Thalaiyasingam Ajanthan<sup>1</sup> Vibhav Vineet<sup>2</sup> Richard Hartley<sup>1</sup>

<sup>1</sup>Australian National University and Australian Centre for Robotic Vision

<sup>2</sup>Microsoft Research, Redmond, USA

<sup>3</sup>Data61, CSIRO, Australia

{zhiwei.xu, thalaiyasingam.ajanthan, richard.hartley}@anu.edu.au

vibhav.vineet@microsoft.com

## Abstract

Although 3D Convolutional Neural Networks (CNNs) are essential for most learning based applications involving dense 3D data, their applicability is limited due to excessive memory and computational requirements. Compressing such networks by pruning therefore becomes highly desirable. However, pruning 3D CNNs is largely unexplored possibly because of the complex nature of typical pruning algorithms that embeds pruning into an iterative optimization paradigm. In this work, we introduce a Resource Aware Neuron Pruning (RANP) algorithm that prunes 3D CNNs at initialization to high sparsity levels. Specifically, the core idea is to obtain an importance score for each neuron based on their sensitivity to the loss function. This neuron importance is then reweighted according to the neuron resource consumption related to FLOPs or memory. We demonstrate the effectiveness of our pruning method on 3D semantic segmentation with widely used 3D-UNets on ShapeNet and BraTS'18 as well as on video classification with MobileNetV2 and I3D on UCF101 dataset. In these experiments, our RANP leads to roughly **50%-95% reduction in FLOPs and 35%-80% reduction in memory** with negligible loss in accuracy compared to the unpruned networks. This significantly reduces the computational resources required to train 3D CNNs. The pruned network obtained by our algorithm can also be easily scaled up and transferred to another dataset for training.

## 1. Introduction

3D image analysis is important in various real-world applications including scene understanding [44, 43], object recognition [16, 14], medical image analysis [5, 42, 20], and video action recognition [18, 33]. Typically, sparse 3D data can be represented using point clouds [38] whereas volumetric representation is required for dense 3D data which arises in domains such as medical imaging [26] and video segmentation and classification [44, 18, 33]. While efficient

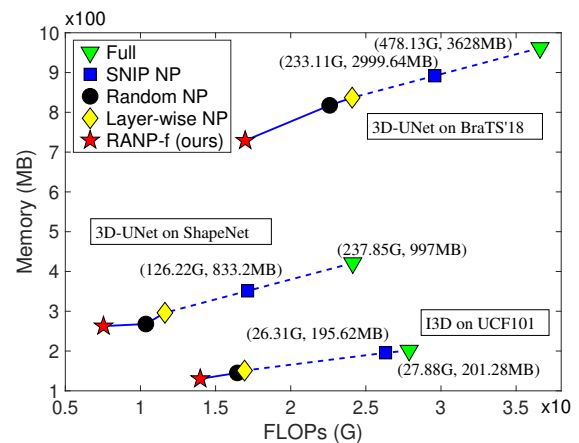


Figure 1: **Bottom-left is the best.** Comparison of neuron pruning methods. “Full” and “SNIP NP” are not drawn by scale but with their FLOPs (G) and memory (MB) values next to the markers. Our RANP-f performs best with large resource reductions while maintaining the accuracy. More details are in Table 2.

neural network architectures can be designed for sparse point cloud data [9, 28], conventional dense 3D Convolutional Neural Network (CNN) is required for volumetric data. Such 3D CNNs are computationally expensive with excessive memory requirements for large-scale 3D tasks. Therefore, it is highly desirable to reduce the memory and FLOPs required to train 3D CNNs while maintaining the accuracy. This will not only enable large-scale applications but also 3D CNN training on resource-limited devices.

Network pruning is a prominent approach to compress a neural network by reducing the number of parameters or the number of neurons in each layer [10, 6, 13, 11]. However, most of the network pruning methods aim at 2D CNNs while pruning 3D CNNs is largely unexplored. This is mainly because pruning is typically targeted at reducing the test-time resource requirements while computational requirements of training time are as large as (if not more

than) the unpruned network. Such pruning schemes are not suitable for 3D CNNs with dense volumetric data where training-time resource requirement is prohibitively large.

In this work, we introduce a Resource<sup>1</sup> Aware Neuron Pruning (RANP) that *prunes 3D CNNs at initialization*. Our method is inspired by, but superior to, SNIP [23] which prunes redundant parameters of a network at initialization and only tests with small scale 2D CNNs for image classification. With the same characteristics of effectively pruning at initialization without requiring large computational resources, RANP yields better-pruned networks compared to SNIP by removing neurons that largely contribute to the high resource requirement. In our experiments on video classification and more challenging 3D semantic segmentation, with minimal accuracy loss, RANP yields 50%-95% reduction in FLOPs and 35%-80% reduction in memory while only 5%-51% reduction in FLOPs and 1%-17% reduction in memory are achieved by SNIP NP.

The main idea of RANP is to prune based on a *neuron importance* criterion analogous to the connection sensitivity in SNIP. Note that, pruning based on such a simple criterion as SNIP has the risk of pruning the whole layer(s) at extreme sparsity levels especially on large networks [22]. Even though an orthogonal initialization that ensures layer-wise dynamical isometry is sufficient to mitigate this issue for parameter pruning on 2D CNNs [22], it is unclear if this could be directly applied to neuron pruning on 3D CNNs. To tackle this and improve pruning, we introduce a *resource aware reweighting scheme* that first balances the mean value of neuron importance in each layer and then reweights the neuron importance based on the resource consumption of each neuron. As evidenced by our experiments, such a reweighting scheme is crucial to obtain large reductions in memory and FLOPs while maintaining high accuracy.

We firstly evaluate our RANP on 3D semantic segmentation on a sparse point-cloud dataset, ShapeNet [38], and a dense medical image dataset, BraTS'18 [26, 1], with widely used 3D-UNets [5]. We also evaluate RANP on video classification using UCF101 with MobileNetV2 [31] and I3D [2]. Our RANP-f significantly outperforms other neuron pruning methods in *resource efficiency* by yielding large reductions in computational resources (**50%-95% FLOPs reduction and 35%-80% memory reduction**) with comparable accuracy to the unpruned network (ref. Fig. 1).

Furthermore, we perform extensive experiments to demonstrate **1) scalability** of RANP by pruning with a small input spatial size and training with a large one, **2) transferability** by pruning using ShapeNet and training on BraTS'18 and vice versa, **3) lightweight** training on a single GPU, and **4) fast** training with increased batch size.

<sup>1</sup>We concretely define “resource” as Floating Point Operations per second (FLOPs) and memory required by one forward pass.

## 2. Related Work

Previous works of network pruning mainly focus on 2D CNNs by parameter pruning [23, 11, 10, 6, 3] and neuron pruning [40, 13, 25, 41, 15, 12]. While a majority of the pruning methods use the traditional prune-retrain scheme with a combined loss function of pruning criteria [11, 13, 41], some pruning at initialization methods is able to reduce computational complexity in training [23, 45, 24, 39, 36]. While very few are for 3D CNNs [27, 46, 4], none of them prune networks at initialization, and thus, none of them effectively reduce the training-time computational and memory requirements of 3D CNNs.

**2D CNN pruning.** *Parameter pruning* merely sparsifies filters for a high learning capability with small models. Han *et al.* [11] adopted an iterative method of removing parameters with values below a threshold. Lee *et al.* [23] recently proposed a single-shot method with connection sensitivity by magnitudes of parameter mask gradients to retain top- $\kappa$  parameters. These filter-sparse methods, however, do not directly yield large speedup and memory reductions.

By contrast, *neuron pruning*, also known as filter pruning or channel pruning, can effectively reduce computational resources. For instance, Li *et al.*[25] used  $l_1$  normalization to remove unimportant filters with connecting features. He *et al.*[13] adopted a LASSO regression to prune network layers with reconstruction in the least square manner. Yu *et al.*[40] proposed a group-wise 2D-filter pruning from each 3D-filter by a learning-based method and a knowledge distillation. Structure learning based MorphNet [8] and SSL [37] aim at pruning activations with structure constraints or regularization. These approaches only reduce the test-time resource requirement while we focus on reducing the resource requirement of large 3D CNNs at training time.

**3D CNN pruning.** To improve the efficiency on 3D CNNs, some works like SSC [9] and OctNet [30] use efficient data structures to reduce the memory requirement for sparse point-cloud data. However, these approaches are not useful for dense data, *e.g.*, MRI images and videos, and the resource requirement remains prohibitively large.

Hence, it is desirable to develop an efficient pruning for 3D CNNs that can handle dense 3D data which is common in real applications. Only very few works are relevant to 3D CNN pruning. Molchanov *et al.*[27] proposed a greedy criteria-based method to reduce resources via back-propagation with a small 3D CNN for hand gesture recognition. Zhang *et al.*[46] used a regularization-based pruning method by assigning regularization to weight groups with  $4\times$  speedup in theory. Recently, Chen *et al.*[4] converted 3D filters into frequency domain to eliminate redundancy in an iterative optimization for convergence. Being a parameter pruning method, this does not lead to large FLOPs and memory reductions, *e.g.*, merely a  $2\times$  speedup compared to our  $28\times$  (ref. Sec. 5.3). In summary, these methods embed

pruning in the iterative network optimization and require extensive resources, which is inefficient for 3D CNNs.

**Pruning at Initialization.** While few works adopted pruning at initialization, some achieved impressive success. SNIP [23] is the first single-shot pruning method that presented a high possibility of pruning networks at initialization with minimal accuracy loss in training, followed by many recent works on single-shot pruning [45, 24, 39, 36]. But none are for 3D CNNs pruning.

In addition to being a parameter pruning approach, the benefits of SNIP was demonstrated only on small-scale datasets [23], such as MNIST and CIFAR-10. Therefore, it is unclear that whether these benefits could be transposed to 3D CNNs applied to large-scale datasets. Our experiments indicate that, while SNIP itself is not capable of yielding large resource reduction on 3D CNNs, our RANP can greatly reduce the computational resources without causing network infeasibility. Furthermore, we show that RANP enjoys strong transferability among datasets and enables fast and lightweight training of large 3D volumetric data segmentation on a single GPU.

### 3. Preliminaries

We first briefly describe the main idea of SNIP [23] which removes redundant parameters prior to training. Given a dataset  $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^S$  with input  $\mathbf{x}_i$  and ground truth  $\mathbf{y}_i$  and the sparsity level  $\kappa$ , the optimization problem associated with SNIP can be written as

$$\begin{aligned} \min_{\mathbf{c}, \mathbf{w}} L(\mathbf{c} \odot \mathbf{w}; \mathcal{D}) &= \min_{\mathbf{c}, \mathbf{w}} \frac{1}{S} \sum_{i=1}^S \ell(\mathbf{c} \odot \mathbf{w}, (\mathbf{x}_i, \mathbf{y}_i)), \quad (1) \\ \text{s.t. } \mathbf{w} &\in \mathbb{R}^m, \quad \mathbf{c} \in \{0, 1\}^m, \quad \|\mathbf{c}\|_0 \leq \kappa, \end{aligned}$$

where  $\mathbf{w}$  is denoted a  $m$ -dimensional vector of parameters,  $\mathbf{c}$  is the corresponding binary mask on the parameters,  $\ell(\cdot)$  is the standard loss function (e.g., cross-entropy loss), and  $\|\cdot\|_0$  denotes  $l_0$  norm. The mask  $c_j \in \{0, 1\}$  for parameter  $w_j$  denotes that the parameter is retained in the compressed model if  $c_j = 1$  and otherwise it is removed. In order to optimize the above problem, they first relax the binary constraint on the masks such that  $\mathbf{c} \in [0, 1]^m$ . Then an importance function for parameter  $w_j$  is calculated by the normalized magnitude of the loss gradient over mask  $c_j$  as

$$s_j = \frac{|g_j(\mathbf{w}; \mathcal{D})|}{\sum_{k=1}^m |g_k(\mathbf{w}; \mathcal{D})|}, \quad \text{where } g_j(\mathbf{w}; \mathcal{D}) = \left. \frac{\partial L(\mathbf{c} \odot \mathbf{w}; \mathcal{D})}{\partial c_j} \right|_{\mathbf{c}=1}. \quad (2)$$

Then, only top- $\kappa$  parameters are retained based on the parameter importance, called connection sensitivity in [23], as defined above. Upon pruning, the retained parameters are trained in the standard way. It is interesting to note that, even though having the mask  $\mathbf{c}$  is easier to explain the intuition, SNIP can be implemented without these additional variables by noting that  $g_j(\mathbf{w}; \mathcal{D}) =$




Figure 2: Flowchart of RANP algorithm. The refining generates a new yet slim network for resource-efficient training.

$(\partial L(\mathbf{w}; \mathcal{D}) / \partial w_j) w_j$  [22]. This method has shown remarkable results in achieving  $> 95\%$  sparsity on 2D image classification tasks with minimal loss of accuracy. Such a parameter pruning method is important, however, it cannot lead to sufficient computation and memory reductions to train a deep 3D CNN on current off-the-shelf graphics hardware. In particular, the sparse weight matrices cannot efficiently reduce memory or FLOPs, and they require specialized sparse matrix implementations for speedup. In contrast, neuron pruning directly translates into practical gains of reducing both memory and FLOPs. This is crucial in 3D CNNs due to their substantially higher resource requirement compared to 2D CNNs.

### 4. Resource Aware NP at Initialization

To explain the proposed RANP, we first extend the SNIP idea to neuron pruning at initialization. Then we discuss a resource aware reweighting strategy to further reduce the computational requirements of the pruned network. The flowchart of our RANP algorithm is shown in Fig. 2.

Before introducing our neuron importance, we first consider a fully-connected feed-forward neural network for the simplicity of notations. Consider weight matrices  $\mathbf{W}^l \in \mathbb{R}^{N_l \times N_{l-1}}$ , biases  $\mathbf{b}^l \in \mathbb{R}^{N_l}$ , pre-activations  $\mathbf{h}^l \in \mathbb{R}^{N_l}$ , and post-activations  $\mathbf{x}^l \in \mathbb{R}^{N_l}$ , for layer  $l \in \mathcal{K} = \{1, \dots, K\}$ . Now the feed-forward dynamics is

$$\mathbf{x}^l = \phi(\mathbf{h}^l), \quad \text{where } \mathbf{h}^l = \mathbf{W}^l \mathbf{x}^{l-1} + \mathbf{b}^l, \quad (3)$$

where the activation function  $\phi: \mathbb{R} \rightarrow \mathbb{R}$  has element-wise nonlinearity and the network input is denoted by  $\mathbf{x}^0$ . Now we introduce binary masks on neurons (i.e., post-activations). The feed-forward dynamics is then modified to include this masking operation as

$$\mathbf{x}^l = \mathbf{c}^l \odot \phi(\mathbf{h}^l), \quad \text{where } \mathbf{c}^l \in \{0, 1\}^{N_l}, \quad \forall l \in \mathcal{K}, \quad (4)$$

where neuron mask  $c_u^l = 1$  indicates neuron  $x_u^l$  is retained and otherwise pruned. Here, neuron pruning can be written as the following optimization problem

$$\begin{aligned} \min_{\mathbf{w}} L(\mathbf{c}, \mathbf{w}; \mathcal{D}) &= \min_{\mathbf{c}, \mathbf{w}} \frac{1}{S} \sum_{i=1}^S \ell(\mathbf{c}, \mathbf{w}; (\mathbf{x}_i, \mathbf{y}_i)), \quad (5) \\ \text{s.t. } \mathbf{w} &\in \mathbb{R}^m, \quad \mathbf{c} \in \{0, 1\}^n, \quad \|\mathbf{c}\|_0 \leq \kappa, \end{aligned}$$

where  $n$  is the total number of neurons and  $\ell(\mathbf{c}, \cdot; \cdot)$  denotes a standard loss function of the feed-forward mapping with

neuron masks  $\mathbf{c}$  defined in Eq. 4. This can be easily extended to convolutional and skip-concatenation operations.

As removing neurons could largely reduce memory and FLOPs compared to merely sparsifying parameters, the core of our approach is benefited by removing redundant neurons from the model. We use an influence function concept developed for parameters to establish neuron importance through the loss function, to locate redundant neurons.

#### 4.1. Neuron Importance

Note that, neuron importance can be derived from the SNIP-based parameter importance discussed in Sec. 3. Another approach is to directly define neuron importance as the normalized magnitude of the neuron mask gradients analogous to parameter importance.

##### Neuron Importance with Parameter Mask Gradients.

The first approach to calculate neuron importance depends on the magnitude of parameter mask gradients, denoted as Magnitude of Parameter Mask Gradients (MPMG). Thus, the importance of neuron  $x_u^l$  is

$$s_u^l = f\left(|g_{u1}^l|, \dots, |g_{uN_{l-1}}^l|\right), \quad (6)$$

where  $g_{uv}^l = \partial L(\mathbf{c} \odot \mathbf{w}; \mathcal{D}) / \partial c_{uv}^l$  with  $c_{uv}^l$  as the mask of parameter  $w_{uv}^l$ . Refer to Eq. 2. Here,  $f(\cdot) : \mathbb{R}^{N_{l-1}} \rightarrow \mathbb{R}$  is a function mapping a set of values to a scalar. We choose  $f(\cdot) = \text{sum}(\cdot)$  with alternatives, *i.e.*, mean and max functions, in Appendix D. Now, we set neuron masks as 1 for neurons with top- $\kappa$  largest neuron importance.

##### Neuron Importance with Neuron Mask Gradients.

Another approach is to directly compute mask gradients on neurons and treat their magnitudes as neuron importance, denoted as Magnitude of Neuron Mask Gradients (MNMG). The neuron importance of  $x_u^l$  is calculated by

$$s_u^l = \left| \frac{\partial L(\mathbf{c}, \mathbf{w}; \mathcal{D})}{\partial c_u^l} \Big|_{\mathbf{c}=1} \right|. \quad (7)$$

Noting that a non-linear activation function  $\phi(\cdot)$  in CNN including but not limited to ReLU can satisfy  $\phi(ch) = c\phi(h), \forall c \geq 0$ . Given such a homogeneous function, the calculation of neuron importance with neuron masks can be derived from parameter mask gradients in the form of

$$\frac{\partial L(\mathbf{c}, \mathbf{w}; \mathcal{D})}{\partial c_u^l} \Big|_{\mathbf{c}=1} = \sum_{v=1}^{N_{l-1}} \frac{\partial L(\mathbf{c} \odot \mathbf{w}; \mathcal{D})}{\partial c_{uv}^l} \Big|_{\mathbf{c}=1}. \quad (8)$$

Details of the influence of such an activation function on neuron importance are provided in Appendix B. These two approaches for neuron importance are in a similar form that while MPMG is by the sum of magnitudes, MNMG is by the magnitude of the sum of parameter mask gradients. It can be implemented directly from parameter gradients.

The neuron importance based on MPMG or MNMG approach can be used to remove redundant neurons. However,

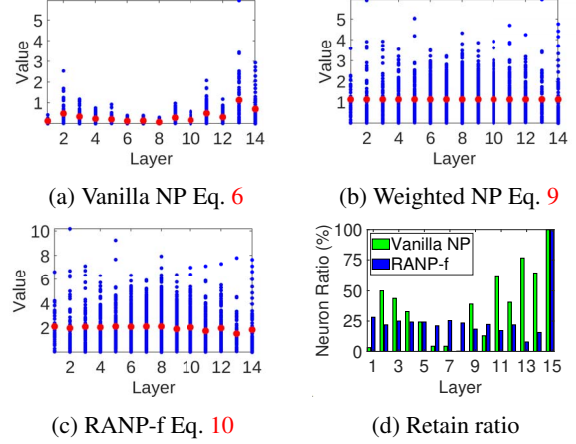


Figure 3: ShapeNet: neuron importance of 3D-UNet becomes balanced and resource-aware from (a) to (c) at neuron sparsity 78.24%. Blue: neuron importance; red: mean values. More illustrations are in Appendix D.

they could lead to an imbalance of sparsity levels of each layer in 3D network architectures. As shown in Table 2, the computational resources required by vanilla neuron pruning are much higher than those by other sparsity enforcing methods, *e.g.*, random neuron pruning and layer-wise neuron pruning. We hypothesize that this is caused by the layer-wise imbalance of neuron importance which unilaterally emphasizes on some specific layer(s) and may lead to network infeasibility by pruning the whole layer(s). This behavior is also observed in [22], and orthogonal initialization is thus recommended to solve the problem for 2D CNN pruning, which however cannot result in balanced neuron importance in our case, see results in Appendix D.

In order to resolve this issue, we propose resource aware neuron pruning (RANP) with reweighted neuron importance, and the details are provided below.

#### 4.2. Resource Aware Reweighting

To tackle the imbalanced neuron importance issue above, we first weight the neuron importance across layers. Weighting neuron importance of  $x_u^l$  can be expressed as

$$\tilde{s}_u^l = \frac{\max_{k=1}^K \bar{s}^k}{\bar{s}^l} s_u^l, \quad \text{where } \bar{s}^k = \frac{1}{N_k} \sum_{u=1}^{N_k} s_u^k, \quad \forall k \in \mathcal{K}. \quad (9)$$

Here,  $\bar{s}^l$  is the mean neuron importance of layer  $l$  and  $\tilde{s}_u^l$  is the updated neuron importance. This helps to achieve the same mean neuron importance in each layer, which largely avoids underestimating neuron importance of specific layer(s) to prevent from pruning the whole layer(s).

To further reduce the memory and FLOPs with minimal accuracy loss, we then reweight the neuron importance  $\tilde{s}_u^l$  by available resource, *i.e.*, memory or FLOPs. This reweighting counts on the addition of weighted neuron importance and the effect of the computational resource, de-

noted as RANP-[m|f], where “m” is for memory and “f” is for FLOPs. We represent the importance of this available resource in layer  $l$  as  $\tau_l$ , refer to Appendix C for details.

The reweighted neuron importance of neuron  $x_u^l$  by following weighted addition variant RANP-[m|f] is

$$\hat{s}_u^l = (1 + \lambda \text{softmax}(-\tau_l)) \tilde{s}_u^l = \left( 1 + \lambda \frac{e^{-\tau_l}}{\sum_{k=1}^K e^{-\tau_k}} \right) \tilde{s}_u^l, \quad (10)$$

where coefficient  $\lambda > 0$  helps to control the effect of resource on neuron importance. This effect represented by softmax constrains the values into a controllable range  $[0, 1]$ , making it easy to determine  $\lambda$  and function a high resource influence with a small resource occupation.

We demonstrate the effect of this reweighting strategy over vanilla pruning in Fig. 3. In more detail, vanilla neuron importance tends to have high values in the last few layers, making it highly possible to remove all neurons of such as the 7th and 8th layers. Weighting the importance in Fig. 3b makes the distribution of importance balanced with the same mean value in each layer. Furthermore, since some neurons have different numbers of input channels, each layer requires different FLOPs and memory. Considering the effect of computational resources on training, we embed them into neuron importance as weights.

In Fig. 3c, the last few layers require larger computational resources than the others, and thus their neurons share lower weights, see the tendency of mean values. Vividly, neuron ratio in Fig. 3d indicates a more balanced distribution by RANP-f than vanilla NP. For instance, very few neurons are retained in the 8th layer by vanilla NP, resulting in low accuracy and low maximum neuron sparsity. With reweighting by RANP-f, however, more neurons can be retained in this layer. Moreover, in Table 2, while weighted NP achieves high accuracy, its computational resource reductions are small. In contrast, RANP-f largely decreases the computational resources with a small accuracy loss.

Then, with reweighted neuron importance by Eq. 10 and  $\tilde{s}_\kappa$  as the  $\kappa$ th reweighted neuron importance in a descending order, the binary mask of neuron  $x_u^l$  can be obtained by

$$c_u^l = 1[\hat{s}_u^l - \tilde{s}_\kappa \geq 0]. \quad (11)$$

As mentioned in Sec. 2, our RANP is more effective in reducing memory and FLOPs than SNIP-based pruning which merely sparsifies parameters but needs high memory required by dense operations in training. RANP can easily remove neurons and all involved input channels at once, leading to huge reductions of input and output channels of the filter. Pseudocode is provided in Appendix A.

## 5. Experiments

We evaluated RANP on 3D-UNets for 3D semantic segmentation and MobileNetV2 and I3D for video classification. Experiments are on Nvidia Tesla P100-SXM2-16GB

GPUs in PyTorch. More results are in Appendix D. Our code is available at <https://github.com/zwxu064/RANP.git>.

### 5.1. Experimental Setup

**3D Datasets.** For 3D semantic segmentation, we adopted the large-scale 3D sparse point-cloud dataset, ShapeNet [38], and dense biomedical MRI sequences, BraTS’18 [26, 1]. *ShapeNet* consists of 50 object part classes, 14007 training samples, and 2874 testing samples. We split it into 6955 training samples and 7052 validation samples as [9] to assign each point/voxel with a part class.

*BraTS’18* includes 210 High Grade Glioma (HGG) and 75 Low Grade Glioma (LGG) cases. Each case has 4 MRI sequences, *i.e.*, T1, T1\_CE, T2, and FLAIR. The task is to detect and segment brain scan images into 3 categories: Enhancing Tumor (ET), Tumor Core (TC), and Whole Tumor (WT). The spatial size is  $240 \times 240 \times 155$  in each dimension. We adopted the splitting strategy of cross-validation in [17] with 228 cases for training and 57 cases for validation.

For video classification, we used video dataset, *UCF101* [34] with 101 action categories and 13320 videos. 2D spatial dimension from images and temporal dimension from frames are cast as dense 3D inputs. Among the 3 official train/test splits, we used split-1 which has 9537 videos for training and 3783 videos for validation.

**3D CNNs.** For 3D semantic segmentation on ShapeNet (sparse data) and BraTS’18 (dense data), we used the standard 15-layer 3D-UNet [5] including 4 encoders, each consists of two “3D convolution + 3D batch normalization + ReLU”, a “3D max pooling”, four decoders, and a confidence module by softmax. It has 14 convolution layers with  $3^3$  kernels and 1 layer with  $1^3$  kernel.

For video classification, we used the popular MobileNetV2 [31, 21] and I3D (with inception as backbone) [2] on UCF101. MobileNetV2 has a linear layer and 52 convolution layers while 18 of them are  $3^3$  kernels and the rest are  $1^3$ . I3D has a linear layer and 57 convolution layers, 19 of which are  $3^3$  kernels, 1 is  $7^3$ , and the rest are  $1^3$ .

**Hyper-parameters in learning.** For *ShapeNet*, we set learning rate as 0.1 with an exponential decay rate  $\gamma = 0.04$  by 100 epochs; batch size is 12 on 2 GPUs; spatial size for pruning and training is  $64^3$  while the spatial size for training is  $128^3$  in Sec. 5.7; optimizer is SGD-Nesterov [35] with weight decay 0.0001 and momentum 0.9.

For *BraTS’18*, learning rate is 0.001, decayed by 0.1 at 150th epoch with 200 epochs; optimizer is Adam[19] with weight decay 0.0001 and AMSGrad[29]; batch size is 2 on 2 GPUs; spatial size for pruning is  $96^3$  and  $128^3$  for training.

For *UCF101*, we adopted similar setup from [21] with learning rate 0.1, decayed by 0.1 at {40, 55, 60, 70}th epoch; optimizer by SGD with weight decay 0.001; batch size 8 on one GPU. Spatial size for pruning and training is  $112^2$  for MobileNetV2 and  $224^2$  for I3D; 16 frames are

Table 1: Vanilla NP by max neuron sparsity. “Metric” is mIoU for ShapeNet, ET for BraTS’18, top-1 for UCF101. “Param” and “Mem” are in MB. MPMG-sum is vanilla NP for large resource reductions and small metric loss.

Dataset (Model)	Manner	Sparsity	Param	GFLOPs	Mem	Metric
ShapeNet (3D-UNet)	Full[5]	0	62.26	237.85	997.00	83.79±0.21
	MNMG-sum	66.93	4.29	100.34	783.14	<b>83.65±0.02</b>
	MPMG-sum	78.24	2.54	<b>55.69</b>	<b>557.32</b>	83.26±0.14
BraTS’18 (3D-UNet)	Full[5]	0	15.57	478.13	3628.00	72.96±0.60
	MNMG-sum	81.32	0.35	<b>73.50</b>	<b>1933.20</b>	64.48±1.10
	MPMG-sum	78.17	0.55	104.50	1936.44	<b>71.94±1.68</b>
UCF101 (MobileNetV2)	Full[31]	0	9.47	0.58	157.47	47.08±0.72
	MNMG-sum	39.89	4.66	<b>0.43</b>	<b>120.01</b>	1.03±0.00 <sup>2</sup>
	MPMG-sum	33.15	6.35	0.55	155.17	<b>46.32±0.79</b>
UCF101 (I3D)	Full[2]	0	47.27	27.88	201.28	51.58±1.86
	MNMG-sum	32.87	20.00	<b>16.03</b>	<b>125.17</b>	49.02±3.33
	MPMG-sum	25.32	29.93	25.76	192.42	<b>51.57±1.46</b>

used for the temporal size. Note that in [21] networks for UCF101 had higher performance since they were pretrained on Kinetics600, while we directly trained on UCF101. A feasible train-from-scratch reference could be [34].

For Eq. 10, we empirically set the coefficient  $\lambda$  as 11 for ShapeNet, 15 for BraTS’18, and 80 for UCF101. Glorot initialization [7] was used for weight initialization. Note that we used orthogonal initialization [32] to handle imbalanced layer-wise neuron importance distribution [22] but obtained lower maximum neuron sparsity.

In addition, loss function and metrics are in Appendix A.

## 5.2. Maximum Neuron Sparsity by Vanilla NP

We selected MPMG-sum and MNMG-sum for vanilla neuron importance for comparison. All neurons of the last convolutional layer are retained for the given classes.

In Table 1, MPMG-sum for ShapeNet achieves the largest neuron sparsity 78.24% by reducing 76.59% FLOPs, 95.92% parameters, and 44.10% memory with 0.53% accuracy loss. Meanwhile, for BraTS’18, MNMG-sum achieves the largest neuron sparsity 81.32% but has up to 8.48% accuracy loss. MPMG-sum, however, has the largest neuron sparsity 78.17% but smaller accuracy loss with decreased 78.14% FLOPs, 96.46% parameters, and 46.63% memory.

Hence, we selected MPMG-sum as vanilla NP considering the trade-off between the maximum neuron sparsity and the accuracy loss. This is applied to all methods related to weighted neuron pruning and RANP in our experiments. Results of mean and max are in Appendix D.

## 5.3. Evaluation of RANP on Pruning Capability

Random NP retains  $\kappa$  neurons with neuron indices randomly shuffled. Layer-wise NP retains neurons using the same retain rate as  $\kappa$  in each layer. For SNIP-based parameter pruning, the parameter masks are post-processed by removing redundant parameters and then making sparse filters

<sup>3</sup>Since 2 layers of the pruned MobileNetV2 by MNMG-sum have only 1 neuron due to the imbalanced layer-wise neuron importance distribution.

dense, which is denoted as SNIP NP. For a fair comparison with SNIP NP, we used the maximum parameter sparsity 98.98% for ShapeNet, 98.88% for BraTS’18, 86.26% for MobileNetV2, and 81.09% for I3D.

**ShapeNet.** Compared with random NP and layer-wise NP in Table 2, the maximum reduced resources by vanilla NP are much less due to the imbalanced layer-wise distribution of neuron importance. Weighted neuron importance by Eq. 9, however, further reduces 18.3% FLOPs and 29.6% memory with 0.14% accuracy loss.

Reweighting by RANP-f and RANP-m further reduces FLOPs and memory on the basis of weighted NP. Here, RANP-f can reduce 96.8% FLOPs, 95.3% parameters, and 73.7% memory over the unpruned networks. Furthermore, with a similar resource in Table 3, RANP achieves  $\sim 0.5\%$  increase in accuracy. Note that a too-large  $\lambda$  can additionally reduce the resources but at the cost of accuracy.

**BraTS’18.** In Table 2, RANP-f achieves 96.5% FLOPs, 95.1% parameters, and 80% memory reductions. It further reduces 18.3% FLOPs and 33.3% memory over vanilla NP while increasing -1.21% ET, 5.11% TC, and 0.77% WT. With a similar resource in Table 3, RANP achieves higher accuracy than random NP and layer-wise NP.

Additionally, Chen *et al.*[4] achieved  $2\times$  speedup on BraTS’18 with 3D-UNet. In comparison, our RANP-f has roughly  $28\times$  speedup, which is theoretically evidenced by the reduced FLOPs from 478.13G to 16.97G in Table 2.

**UCF101.** In Table 2, for MobileNetV2, RANP-f reduces 55.2% FLOPs, 49% parameters, and 44.1% memory with around 1% accuracy loss. Meanwhile, for I3D, it reduces 49.9% FLOPs, 43.5% parameters, and 35.3% memory with around 2% accuracy increase. The RANP-based methods can reduce much more resources than other methods.

## 5.4. Resources and Accuracy with Neuron Sparsity

Here, we further studied the tendencies of resources and accuracy with an increasing neuron sparsity level from 0 to the maximum one with network feasibility.

**Resource Reductions.** In Figs. 4a-4d, RANP, marked with (w), achieves much larger FLOPs and memory reductions than vanilla NP, marked with (w/o), due to the balanced distribution of neuron importance by reweighting.

Specifically, for *ShapeNet*, RANP prunes up to 98.57% neurons while only up to 78.24% by vanilla NP in Fig. 4a. For *BraTS’18*, RANP can prune up to 96.24% neurons while only up to 78.17% neurons can be pruned by vanilla NP in Fig. 4b. For *UCF101*, RANP can prune up to 80.83% neurons compared to 33.15% on MobileNetV2 in Fig. 4c, and 85.3% neurons compared to 25.32% on I3D in Fig. 4d.

**Accuracy with Pruning Sparsity.** For *ShapeNet* in Fig. 4e, the 23-layer 3D-UNet achieves a higher mIoU than the 15-layer one. Extremely, when pruned with the maximum neuron sparsity 97.99%, it can achieve 78.10% mIoU. With the maximum neuron sparsity 98.57%, however, the

Table 2: Evaluation of neuron pruning capability. All models are **trained from scratch** for 100 epochs on ShapeNet and UCF101, 200 on BraTS’18. Metrics are calculated by the last 5 epochs. “sparsity” is max parameter sparsity for SNIP NP and max neuron sparsity for others. Among the neuron pruning methods, we marked bold **the best** and underlined **the second best**. “↓” denotes reduction in %. Overall, our RANP-f performs best with large reductions of **main resource consumption** (GFLOPs and memory) with negligible accuracy loss.

Dataset	Model	Manner	Sparsity(%)	Param(MB)	GFLOPs	Memory(MB)	Metrics(%)			
ShapeNet [38]	3D-UNet	ours	Full[5]	0	62.26	237.85	997.00	mIoU		
			SNIP[23] NP	98.98	5.31 (91.5↓)	126.22 (46.9↓)	833.20 (16.4↓)	<b>83.70±0.20</b>		
			Random NP		3.05 (95.1↓)	10.36 (95.6↓)	267.95 (73.1↓)	82.90±0.19		
			Layer-wise NP		2.99 (95.2↓)	11.63 (95.1↓)	296.22 (70.3↓)	83.25±0.14		
			Vanilla NP	78.24	2.54 (95.9↓)	55.69 (76.6↓)	557.32 (44.1↓)	<u>83.26±0.14</u>		
			Weighted NP		2.97 (95.2↓)	12.06 (94.9↓)	301.56 (69.8↓)	83.12±0.09		
			RANP-m		3.39 (94.6↓)	<b>6.68 (97.2↓)</b>	<b>214.95 (78.4↓)</b>	82.35±0.24		
			RANP-f		2.94 (95.3↓)	<u>7.54 (96.8↓)</u>	<u>262.66 (73.7↓)</u>	83.07±0.22		
BraTS’18 [26, 1]	3D-UNet	ours	Full[5]	0	15.57	478.13	3628.00	ET	TC	WT
			SNIP[23] NP	98.88	1.09 (93.0↓)	233.11 (51.2↓)	2999.64 (17.3↓)	72.96±0.60	73.51±1.54	86.79±0.35
			Random NP		0.75 (95.2↓)	22.59 (95.3↓)	817.59 (77.5↓)	<b>73.33±1.89</b>	71.98±2.15	<b>86.44±0.39</b>
			Layer-wise NP		0.75 (95.2↓)	24.09 (95.0↓)	836.88 (77.0↓)	67.27±0.99	71.62±1.20	74.16±1.33
			Vanilla NP	78.17	0.55 (96.5↓)	104.50 (78.1↓)	1936.44 (46.6↓)	69.74±1.33	71.49±1.62	<u>86.38±0.39</u>
			Weighted NP		0.79 (95.0↓)	22.40 (95.3↓)	860.64 (76.3↓)	<u>71.94±1.68</u>	69.39±2.29	84.68±0.78
			RANP-m		0.87 (94.4↓)	<b>13.47 (97.2↓)</b>	<b>506.97 (86.0↓)</b>	71.50±0.63	<b>75.05±1.19</b>	84.05±0.65
			RANP-f		0.76 (95.1↓)	<u>16.97 (96.5↓)</u>	<u>729.11 (80.0↓)</u>	66.70±2.94	62.99±2.38	82.90±0.41
UCF101 [34]	MobileNetV2	ours	Full[31]	0	9.47	0.58	157.47	Top-1		Top-5
			SNIP[23] NP	86.26	3.67 (61.3↓)	0.54 ( 6.9↓)	155.35 ( 1.3↓)	47.08±0.72	76.68±0.50	
			Random NP		4.58 (51.6↓)	0.34 (41.4↓)	106.68 (32.3↓)	45.78±0.04	75.08±0.17	
			Layer-wise NP		4.56 (51.8↓)	0.33 (43.1↓)	106.92 (32.1↓)	44.74±0.36	74.69±0.58	
			Vanilla NP	33.15	6.35 (32.9↓)	0.55 ( 5.2↓)	155.17 ( 1.5↓)	44.90±0.36	75.54±0.34	
			Weighted NP		4.82 (49.1↓)	0.30 (48.3↓)	100.33 (36.3↓)	<b>46.32±0.79</b>	75.42±0.60	
			RANP-m		4.87 (48.6↓)	<u>0.27 (53.4↓)</u>	<b>84.51 (46.3↓)</b>	45.11±0.41	75.53±0.37	
			RANP-f		4.83 (49.0↓)	<b>0.26 (55.2↓)</b>	<u>88.01 (44.1↓)</u>	45.87±0.41	<b>75.75±0.30</b>	
	I3D	ours	Full[2]	0	47.27	27.88	201.28	51.58±1.86	77.35±0.63	
			SNIP[23] NP	81.09	30.06 (36.4↓)	26.31 ( 5.6↓)	195.62 ( 2.8↓)	52.38±3.55	78.32±3.24	
			Random NP		26.36 (44.2↓)	16.45 (41.0↓)	145.07 (27.9↓)	52.42±2.52	<u>79.05±2.06</u>	
			Layer-wise NP		26.67 (43.6↓)	16.93 (39.3↓)	150.95 (25.0↓)	52.77±1.99	78.41±1.07	
			Vanilla NP	25.32	29.93 (36.7↓)	25.76 ( 7.6↓)	192.42 ( 4.4↓)	51.57±1.46	78.07±1.34	
			Weighted NP		26.57 (43.8↓)	15.56 (44.2↓)	142.57 (29.2↓)	<u>54.09±0.82</u>	79.26±0.61	
			RANP-m		26.75 (43.4↓)	<u>14.08 (49.5↓)</u>	<u>130.44 (35.2↓)</u>	52.11±3.05	77.54±2.64	
			RANP-f		26.69 (43.5↓)	<b>13.98 (49.9↓)</b>	<b>130.22 (35.3↓)</b>	<b>54.27±2.88</b>	<b>79.27±2.13</b>	

Table 3: In addition to Table 2, with similar GFLOPs or memory on 3D-UNets, our RANP-f achieves the highest accuracy.

Manner	ShapeNet					BraTS’18						
	Sparsity	Param	GFLOPs	Mem	mIoU	Sparsity	Param	GFLOPs	Mem	ET	TC	WT
Random NP	81.01	2.27	~7.54	253.12	82.66±0.23	81.08	0.56	~16.97	685.77	61.09±1.87	68.94±2.44	78.89±2.47
Layer-wise NP	82.82	1.84	~7.54	255.67	82.82±0.26	83.50	0.46	~16.97	700.64	<u>70.50±0.63</u>	<u>74.27±0.95</u>	<u>83.63±0.92</u>
Random NP	78.83	2.87	9.57	~262.66	<u>82.86±0.45</u>	80.90	0.57	17.95	~729.11	68.45±1.11	70.67±1.21	75.02±0.79
Layer-wise NP	82.81	1.94	8.14	~262.66	82.52±0.13	82.45	0.51	17.31	~729.11	70.45±1.03	69.27±1.95	82.42±0.68
RANP-f(ours)	78.24	2.94	7.54	262.66	<b>83.07±0.22</b>	78.17	0.76	16.97	729.11	<b>70.73±0.66</b>	<b>74.50±1.05</b>	<b>85.45±1.06</b>

15-layer 3D-UNet achieves only 61.42%.

For *BraTS’18* in Fig. 4f, the 23-layer 3D-UNet does not always outperform the 15-layer one and has a larger fluctuation which could be caused by the limited training samples. Nevertheless, even in the extreme case, the 23-layer 3D-UNet has small accuracy loss. Clearly, RANP makes it feasible to use deeper 3D-UNets without the memory issue.

For *UCF101* in Figs. 4g-4h, RANP-f achieves <3% accuracy loss at 70% neuron sparsity, indicating its effectiveness of greatly reducing resources with small accuracy loss.

### 5.5. Transferability with Interactive Model

In this experiment, we trained on ShapeNet with a transferred 3D-UNet by RANP on BraTS’18 with 80% neuron

Table 4: Transfer learning by 23-layer 3D-UNets interactively pruned and trained between ShapeNet and BraTS’18. Accuracy loss from RANP-f to T-RANP-f is negligible. “T”: transferred.

Manner	ShapeNet mIoU(%)	BraTS’18		
		ET(%)	TC(%)	WT(%)
Full[5]	<b>84.27±0.21</b>	<b>74.04±1.45</b>	<b>75.11±2.43</b>	<u>84.49±0.74</u>
RANP-f(ours)	<u>83.86±0.15</u>	71.13±1.43	72.40±1.48	83.32±0.62
T-RANP-f(ours)	83.25±0.17	<u>72.74±0.69</u>	<u>73.25±1.69</u>	<b>85.22±0.57</b>

sparsity. Interactively, with the same neuron sparsity, a transferred 3D-UNet by RANP on ShapeNet was applied to train on BraTS’18. Results in Table 4 demonstrate that training with transferred models crossing different datasets can largely maintain high or higher accuracy.

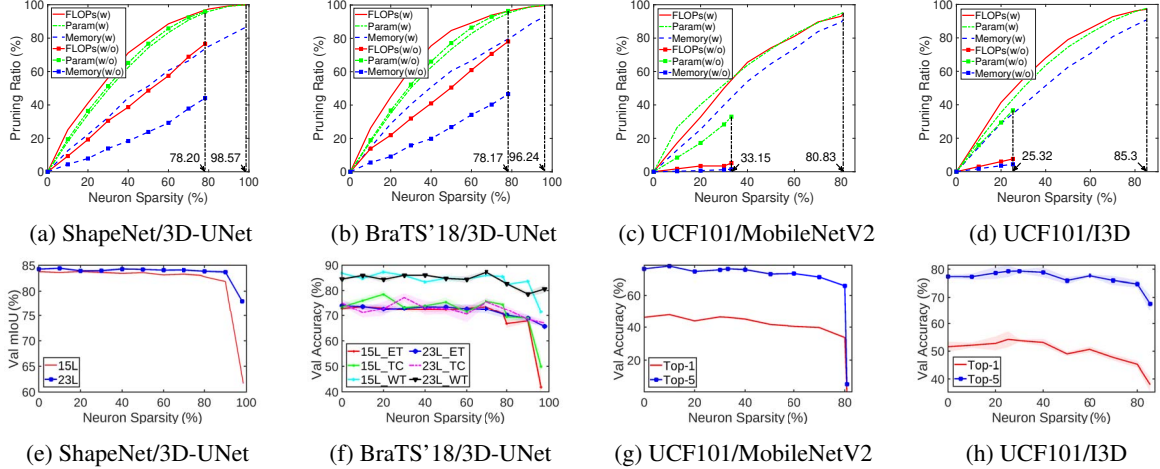


Figure 4: With minimal accuracy loss, more resources are reduced with (w) reweighting by RANP-f than without (w/o) by vanilla NP. (a)-(d) are resources reductions (w) and (w/o) reweighting; (e)-(h) are accuracy and sparsity. Best view in color.

Table 5: ShapeNet: a deeper 23-layer 3D-UNet is achievable on a single GPU with 80% neuron pruning.

Manner	Layer	Batch	GPU(s)	Sparsity(%)	mIoU(%)
Full	15	12	2	0	83.79±0.21
Full	23	12	2	0	<u>84.27±0.21</u>
RANP-f(ours)	23	12	<b>1</b>	80	<b>84.34±0.21</b>

## 5.6. Lightweight Training on a Single GPU

RANP with high neuron sparsity makes it feasible to train with large data size on a single GPU due to the largely reduced resources. We trained on ShapeNet with the same batch size 12 and spatial size  $64^3$  in Sec. 5.1 using a 23-layer 3D-UNet with 80% neuron sparsity on a single GPU. With this setup, RANP-f reduces  $\sim 35\times$  GFLOPs (from 259.59 to 7.39) and  $\sim 3.9\times$  memory (from 1005.96MB to 255.57MB), making it feasible to train on a single GPU instead of 2 GPUs. It achieves a higher mIoU,  $84.34\pm 0.21\%$ , than the 15-layer and 23-layer full 3D-UNets in Table 5.

The accuracy increase is due to the enlarged batch size on each GPU. With limited memory, however, training a 23-layer full 3D-UNet on a single GPU is infeasible.

## 5.7. Fast Training with Increased Batch Size

Here, we used the largest spatial size  $128^3$  of one sample on a single GPU and then extended it to RANP with increased batch size from 1 to 4 to fully fill GPU capacity. The initial learning rate was reduced from 0.1 to 0.01 due to the batch size decreased from 12 in Table 5. This is to avoid an immediate increase in training loss right after 1st epoch due to the unsuitably large learning space.

In Fig. 5a, RANP-f enables increased batch size 4 and achieves a faster loss convergence than the full network. In Fig. 5c, the full network executed 6 epochs while RANP-f reached 26 epochs. Vividly shown by training time in Figs. 5b and 5d, RANP-f has much lower loss and higher accuracy than the full one. This greatly indicates the practical advantage of RANP on fastening training convergence.




Figure 5: ShapeNet: a faster convergence on a single GPU with 23-layer 3D-UNet and increased batch size due to the largely reduced resources by RANP-f. Batch size is 1 for “Full” and 4 for “RANP-f”. Experiments run for 40 hours.

## 6. Conclusion

In this paper, we propose an effective resource aware neuron pruning method, RANP, for 3D CNNs. RANP prunes a network at initialization by greatly reducing resources with negligible loss of accuracy. Its resource aware reweighting scheme balances the neuron importance distribution in each layer and enhances the pruning capability of removing a high ratio, say 80% on 3D-UNet, of neurons with minimal accuracy loss. This advantage enables training deep 3D CNNs with a large batch size to improve accuracy and achieving lightweight training on one GPU.

Our experiments on 3D semantic segmentation using ShapeNet and BraTS’18 and video classification using UCF101 demonstrate the effectiveness of RANP by pruning 70%-80% neurons with minimal loss of accuracy. Moreover, the transferred models pruned on a dataset and trained on another one are succeeded in maintaining high accuracy, indicating the high transferability of RANP. Meanwhile, the largely reduced computational resources enable lightweight and fast training on one GPU with increased batch size.

## Acknowledgement

We would like to thank Ondrej Miksik for valuable discussions. This work is supported by the Australian Centre for Robotic Vision (CE140100016) and Data61, CSIRO.



## References

- [1] S. Bakas, H. Akbari, A. Sotiras, M. Bilello, M. Rozycki, J. Kirby, J. Freymann, K. Farahani, and C. Davatzikos. Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features. *Nature Scientific Data*, 2017.
- [2] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the Kinetics dataset. *CVPR*, 2017.
- [3] C. Chen, F. Tung, N. Vedula, and G. Mori. Constraint-aware deep neural network compression. *ECCV*, 2018.
- [4] H. Chen, Y. Wang, H. Shu, Y. Tang, C. Xu, B. Shi, C. Xu, Q. Tian, and C. Xu. Frequency domain compact 3D convolutional neural networks. *CVPR*, 2020.
- [5] O. Cicek, A. Abdulkadir, S. Lienkamp, T. Brox, and O. Ronneberger. 3D U-Net: Learning dense volumetric segmentation from sparse annotation. *MICCAI*, 2016.
- [6] X. Dong, S. Chen, and S. Pan. Learning to prune deep neural networks via layer-wise optimal brain surgeon. *NeurIPS*, 2017.
- [7] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2010.
- [8] A. Gordon, E. Eban, O. Nachum, B. Chen, H. Wu, T. Yang, and E. Choi. Morphnet: Fast & simple resource-constrained structure learning of deep networks. *CVPR*, 2018.
- [9] B. Graham, M. Engelcke, and L. Maaten. 3D semantic segmentation with submanifold sparse convolutional networks. *CVPR*, 2018.
- [10] Y. Guo, A. Yao, and Y. Chen. Dynamic network surgery for efficient dnns. *NeurIPS*, 2016.
- [11] S. Han, J. Pool, J. Tran, and W. Dally. Learning both weights and connections for efficient neural network. *NeurIPS*, 2015.
- [12] Y. He, J. Lin, Z. Liu, H. Wang, L. Li, and S. Han. Amc: Automl for model compression and acceleration on mobile devices. *ECCV*, 2018.
- [13] Y. He, X. Zhang, and J. Sun. Channel pruning for accelerating very deep neural networks. *ICCV*, 2017.
- [14] R. Hou, C. Chen, R. Sukthankar, and M. Shah. An efficient 3D CNN for action/object segmentation in video. *BMVC*, 2019.
- [15] Z. Huang and N. Wang. Data-driven sparse structure selection for deep neural networks. *ECCV*, 2018.
- [16] S. Ji, W. Xu, M. Yang, and K. Yu. 3D convolutional neural networks for human action recognition. *TPAMI*, 2013.
- [17] P. Kao, T. Ngo, A. Zhang, J. Chen, and B. Manjunath. Brain tumor segmentation and tractographic feature extraction from structural MR images for overall survival prediction. *Workshop on MICCAI*, 2018.
- [18] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. *CVPR*, 2014.
- [19] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *ICLR*, 2015.
- [20] J. Kleesiek, G. Urban, A. Hubert, D. Schwarz, K. Hein, M. Bendszus, and A. Biller. Deep MRI brain extraction: A 3D convolutional neural network for skull stripping. *NeuroImage*, 2016.
- [21] O. Kopuklu, N. Kose, A. Gunduz, and G. Rigoll. Resource efficient 3D convolutional neural networks. *ICCVW*, 2019.
- [22] N. Lee, T. Ajanthan, S. Gould, and P. Torr. A signal propagation perspective for pruning neural networks at initialization. *ICLR*, 2020.
- [23] N. Lee, T. Ajanthan, and P. Torr. SNIP: Single-shot network pruning based on connection sensitivity. *ICLR*, 2019.
- [24] C. Li, Z. Wang, X. Wang, and H. Qi. Single-shot channel pruning based on alternating direction method of multipliers. *arXiv:1902.06382*, 2019.
- [25] H. Li, A. Kadav, I. D. H. Samet, and H. Graf. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710*, 2016.
- [26] B. Menze, A. Jakab, and S. B. *et al.* The multimodal brain tumor image segmentation benchmark (brats). *IEEE Transactions on Medical Imaging*, 2015.
- [27] P. Molchanov, S. Tyree, T. Karras, T. Aila, and J. Kautz. Pruning convolutional neural networks for resource efficient inference. *ICLR*, 2017.
- [28] C. Qi, H. Su, K. Mo, and L. Guibas. Pointnet: Deep learning on point sets for 3D classification and segmentation. *CVPR*, 2017.
- [29] S. Reddi, S. Kale, and S. Kumar. On the convergence of adam and beyond. *ICLR*, 2018.
- [30] G. Riegler, A. Ulusoy, and A. Geiger. OctNet: Learning deep 3D representations at high resolutions. *CVPR*, 2017.
- [31] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. *CVPR*, 2018.
- [32] A. Saxe, J. McClelland, and S. Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *ICLR*, 2014.
- [33] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. *NeurIPS*, 2014.
- [34] K. Soomro, A. Zamir, and M. Shah. UCF101: A dataset of 101 human action classes from videos in the wild. *CRCV-Techinal Report*, 2012.
- [35] I. Sutskever, J. Martens, G. Dahl, and G. Hinton. On the importance of initialization and momentum in deep learning. *ICML*, 2013.
- [36] C. Wang, G. Zhang, and R. Grosse. Picking winning tickets before training by preserving gradient flow. *ICLR*, 2020.
- [37] W. Wen, C. Wu, Y. Wang, Y. Chen, and H. Li. Learning structured sparsity in deep neural networks. *NeurIPS*, 2016.
- [38] L. Yi, V. Kim, D. Ceylan, I. Shen, M. Yan, H. Su, C. Lu, Q. Huang, A. Sheffer, and L. Guibas. A scalable active framework for region annotation in 3D shape collections. *SIGGRAPH Asia*, 2016.
- [39] J. Yu and T. Huang. Autoslim: Towards one-shot architecture search for channel numbers. *arXiv:1903.11728*, 2019.
- [40] N. Yu, S. Qiu, X. Hu, and J. Li. Accelerating convolutional neural networks by group-wise 2D-filter pruning. *IJCNN*, 2017.
- [41] R. Yu, A. Li, C. Chen, J. Lai, V. Morariu, X. Han, M. Gao, C. Lin, and L. Davis. NISP: Pruning networks using neuron importance score propagation. *CVPR*, 2018.

- [42] F. Zanjani, D. Moin, B. Verheij, F. Claessen, T. Cherici, T. Tan, and P. With. Deep learning approach to semantic segmentation in 3D point cloud intra-oral scans of teeth. *Proceedings of Machine Learning Research*, 2019.
- [43] C. Zhang, W. Luo, and R. Urtasun. Efficient convolutions for real-time semantic segmentation of 3D point cloud. *3DV*, 2018.
- [44] H. Zhang, K. Jiang, Y. Zhang, Q. Li, C. Xia, and X. Chen. Discriminative feature learning for video semantic segmentation. *International Conference on Virtual Reality and Visualization*, 2014.
- [45] M. Zhang and B. Stadie. One-shot pruning of recurrent neural networks by jacobian spectrum evaluation. *arXiv:1912.00120*, 2019.
- [46] Y. Zhang, H. Wang, Y. Luo, L. Yu, H. Hu, H. Shan, and T. Quek. Three dimensional convolutional neural network pruning with regularization-based method. *ICIP*, 2019.