

# Structural Pruning via Spatial-aware Information Redundancy for Semantic Segmentation

Dongyue Wu<sup>1</sup>, Zilin Guo<sup>1</sup>, Li Yu<sup>2</sup>, Nong Sang<sup>1</sup>, Changxin Gao<sup>1\*</sup>

<sup>1</sup>National Key Laboratory of Multispectral Information Intelligent Processing Technology, School of Artificial Intelligence and Automation, Huazhong University of Science and Technology

<sup>2</sup>School of Electronic Information and Communications, Huazhong University of Science and Technology  
{dongyue\_wu, zilin\_guo, hustlyu, nsang, cgao}@hust.edu.cn

## Abstract

In recent years, semantic segmentation has flourished in various applications. However, the high computational cost remains a significant challenge that hinders its further adoption. The filter pruning method for structured network slimming offers a direct and effective solution for the reduction of segmentation networks. Nevertheless, we argue that most existing pruning methods, originally designed for image classification, overlook the fact that segmentation is a location-sensitive task, which consequently leads to their suboptimal performance when applied to segmentation networks. To address this issue, this paper proposes a novel approach, denoted as Spatial-aware Information Redundancy Filter Pruning (SIRFP), which aims to reduce feature redundancy between channels. First, we formulate the pruning process as a maximum edge weight clique problem (MEWCP) in graph theory, thereby minimizing the redundancy among the remaining features after pruning. Within this framework, we introduce a spatial-aware redundancy metric based on feature maps, thus endowing the pruning process with location sensitivity to better adapt to pruning segmentation networks. Additionally, based on the MEWCP, we propose a low computational complexity greedy strategy to solve this NP-hard problem, making it feasible and efficient for structured pruning. To validate the effectiveness of our method, we conducted extensive comparative experiments on various challenging datasets. The results demonstrate the superior performance of SIRFP for semantic segmentation tasks.

**Code** — <https://github.com/dyw98/SIRFP.git>

## Introduction

Semantic segmentation is a classic and challenging task in the field of computer vision. Its goal is to assign category labels to each pixel in an image. Recently, with the advent of Fully Convolutional Networks (FCN) (Long, Shelhamer, and Darrell 2015), numerous deep learning-based methods (Chen et al. 2017a; Ding et al. 2018; Ronneberger, Fischer, and Brox 2015; Zhao et al. 2017; Pang et al. 2019; Chen et al. 2018) have been proposed and applied to semantic segmentation. These methods often involve networks with a large number of parameters and high computational

\*Corresponding author.

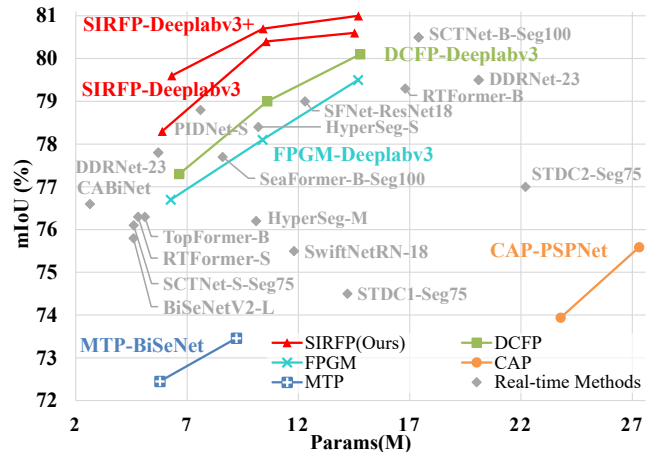


Figure 1: Params vs mIOU on Cityscapes dataset. The proposed Spatial-aware Information Redundancy Filter Pruning (SIRFP) method achieves an outstanding trade-off between parameter efficiency and accuracy. The pruning methods are connected with lines while real-time methods are colored in grey. Best viewed in color.

costs, which enable them to extract high-quality features for pixel representation. However, the intimidating computational cost also hinders their deployment on resource-constrained devices widely used on autonomous vehicles, mobile phones, wearable devices, and robots. Therefore, there is an urgent need to develop efficient neural network compression methods for semantic segmentation.

Currently, network compression methods can be broadly categorized into three types: parameter quantization (Han, Mao, and Dally 2015; Xu et al. 2018; Esser et al. 2019), low-rank approximation (Yang et al. 2019; Yu et al. 2017), and network pruning (Luo, Wu, and Lin 2017; Lin et al. 2020b; Molchanov et al. 2019). Network pruning can be further divided into weight pruning (Han et al. 2015; Carreira-Perpinán and Idelbayev 2018) and filter pruning (Zhuang et al. 2018; Molchanov et al. 2019; Li et al. 2016). Weight pruning eliminates individual weight elements to achieve unstructured sparsity, requiring specialized libraries to accelerate sparse matrix computations. In contrast, filter prun-

ing directly prunes redundant channels (or filters). The pruned network retains the same structure as the original, differing only in the number of channels. Thus, filter pruning can achieve network compression using standard hardware without specially designed libraries. This simplicity has made filter pruning a widely adopted approach, provoking extensive research in this area.

However, most pruning methods are primarily focused on image recognition tasks. Although these methods achieve promising performance in image classification tasks, they fail to account for the differences between image recognition and semantic segmentation. Classification focuses on identifying distinguishing features, whereas segmentation further requires precise object localization within images. Rich spatial information is crucial for high-quality segmentation, but it plays a lesser role in classification. Consequently, semantic segmentation models exhibit different channel redundancy characteristics compared to classification models. Identifying redundancy in segmentation requires evaluating not only the similarity of semantic information but also the spatial redundancy across channels. The spatial-agnostic importance metric adopted in existing filter pruning methods limits their effectiveness on semantic segmentation models, leading to a suboptimal trade-off between compression ratios and accuracy performance.

To address these issues, we propose a spatially-aware information redundancy filter pruning (SIRFP) method for semantic segmentation. To measure the importance of each channel, we first introduce a spatially-aware redundancy metric, statistically calculating the redundancy based on the non-pooled, high-resolution, and spatially informative features during training using the exponential moving average (EMA) mechanism. With this spatially-aware redundancy metric, we rethink the optimization objective of structured pruning from the graph theory perspective, finding it analogous to the classical *Maximum Edge Weight Clique Problem* (MEWCP) (Späth 1985). This problem aims to find the complete subgraph with the maximum sum of all its edge weights of a given cardinality (the number of nodes). From this perspective, we construct the graph of filters: each node represents a filter, and the edge weights are negatively correlated with the mutual information redundancy between filters. With this constructed graph, the filter pruning optimization goal that minimizes the mutual information redundancy between channels in the pruned network is exactly equivalent to MEWCP. Unfortunately, MEWCP is an NP-hard problem, which could not be directly solved in polynomial time. Although various heuristic methods (Späth 1985; Pullan 2008) and exact methods (Hosseinian, Fontes, and Butenko 2020) have been proposed for MEWCP, they still cost unacceptable time in network pruning scenarios where the number of filter nodes is much larger (more than 1024) than the cardinality of graphs (less than 500) that they typically focus on. To address this challenge, we propose an efficient greedy algorithm that significantly reduces the computational complexity and time cost, providing a practical solution to our revised pruning optimization problem. Extensive experiments on popular semantic segmentation datasets and models demonstrate that our method achieves

outstanding results. For example, Figure 1 shows that the proposed SIRFP achieves the state-of-the-art trade-off between the number of parameters and mIoU performance on the Cityscapes dataset compared with both pruning methods and recent real-time segmentation methods. Our method is also evaluated on object detection and image classification tasks, showing decent generalization ability. In summary, the contributions of this paper are as follows:

1. We highlight that current channel pruning methods overlook the spatial redundancy in semantic segmentation models, leading to limited performance and compression ratios.
2. To address this issue, we propose a spatially-aware redundancy metric that directly calculates information redundancy based on high-resolution feature maps that retain spatial location information, using a sliding average during network training to avoid additional inference and feature storage post-training.
3. We propose to minimize spatial-aware redundancy between the left channels in the pruned model and frame the optimization objective as the maximum weight clique problem (MEWCP). A fast and efficient greedy pruning process is presented to solve this NP-hard problem.

## Related Work

### Semantic Segmentation

Since FCN (Long, Shelhamer, and Darrell 2015) introduced a simple and effective solution for semantic segmentation, numerous deep learning methods have continuously pushed the performance boundaries in this field. Compared with classification methods, segmentation approaches are dedicated to aggregating more contextual information and preserving rich spatial information. For example, PSPNet (Zhao et al. 2017) and the DeepLab series (Chen et al. 2017a,b, 2014, 2018) maintain a rather large resolution of feature map for the final prediction. Recent transformer-based methods like ViT (Dosovitskiy et al. 2020) and Swin (Liu et al. 2021b) also enhance spatial information by adopting positional encoding to improve the performance.

In the field of real-time segmentation, features with abundant spatial details are even more important given that the computational cost is under constraints. However, they manage to preserve rich spatial information by adopting light-weighted multi-branch networks. For instance, the BiSeNet (Yu et al. 2018, 2021) series introduces a two-branch architecture to get rich spatial details and high-level semantic information, respectively. STDC (Fan et al. 2021) adopts an additional spatial guidance branch to enhance spatial information in the extracted features. Recently, DDRNet (Pan et al. 2022), RTFormer (Wang et al. 2022), and SeaFormer (Wan et al. 2023) employ a feature-sharing framework which divides spatial and semantic features at the deep stages. All these methods demonstrate the importance of the spatial distribution information for segmentation.

### Network Pruning

Network Pruning is a popular and effective technique for network compression, which aims to eliminate redundant

parameters in the neural network. Opposed to unstructured pruning (weight pruning), structural pruning removes parameters following certain structured sparsity modes. Filter pruning, which removes the entire filters, is one of the most popular techniques to get an efficient model without the support of specially designed hardware or libraries. The key is how to evaluate the importance of each filter and the corresponding channel. Most of them (He et al. 2019; Liu et al. 2017; Fang et al. 2023; Luo, Wu, and Lin 2017) resort to the property of filter weights, while there are also methods (Sui et al. 2021; Lin et al. 2020a) analyze the characteristic of the corresponding feature maps. As most of the existing pruning methods focus on image classification, the spatial distribution in the feature map, which is of significance for location-sensitive tasks such as segmentation and detection, hardly attracts the attention of researchers.

Some recent works are designed specifically for semantic segmentation. CAP (He et al. 2021) first proposes to quantify the contextual information using pooled features. The quantified context is utilized as a regularization term. It penalizes the scaling factors of channels with less contextual information and encourages those with more contextual information. Chen et al. (Chen, Zhang, and Wang 2022) employ a multi-task channel pruning framework where the importance of each channel is simultaneously determined by both the classification and segmentation tasks. DCFP (Wang et al. 2024) pointed out that the long-tail distribution of samples hinders the performance of pruned networks. Thus, a gradient-based importance criterion along with the powerful Geometric-Semantic Re-balanced Loss (GSRL) is proposed to avoid performance loss and boost the learning on those tail classes during fine-tuning. These methods significantly improve the efficiency-accuracy trade-off of pruned segmentation models, but still ignore the great disparity between image classification and semantic segmentation: segmentation requires abundant spatial information to locate objects while classification only needs class-related information for categorizing. Therefore, without rich spatial information, the utilization of only parameters of filters (or channels) or the pooled features prevents them from achieving superb segmentation performance of pruned models.

## Methodology

In this section, we first introduce the optimization problem of most pruning methods. Then, our motivation to address the network pruning issue based on the graph theory is elaborated. After that, we introduce the details of the optimization problem of structural pruning as maximum edge weight clique problem (MEWCP), the proposed spatial-aware information redundancy metric, and our efficient heuristic greedy pruning (EHGP) method to efficiently address the NP-hard MEWCP.

### Preliminary

For clarity, we first define the denotations. Suppose a neural network contains  $L$  layers. The set of network parameters  $\{\mathbf{W}^1, \dots, \mathbf{W}^L\}$  contains all the learnable weights.  $\mathbf{W}^l \in \mathcal{R}^{C^l \times C^{l-1}}$  denotes the weight matrix of  $l$ -th convolu-

tion or linear layer, where  $C^l$  and  $C^{l-1}$  denoting the number of output and input channels, respectively. The  $i$ -th column in  $\mathbf{W}^l$ , namely the weight vector of the  $i$ -th output channel, in the  $l$ -th layer is represented by  $\mathbf{w}_{i,\cdot}^l$ .

For the  $l$ -th layer with channel sparsity  $\kappa^l$ , channel pruning methods identify and retain the set of channels  $\mathcal{T}^l$  containing the most important information to form an efficient light-weighted pruned model. The optimization problem can be formulated as follows:

$$\max_{\mathbf{m}^l} \sum_{i=1}^{C^l} m_i^l \Phi(l, i), \quad s.t. \quad \sum_{j=1}^{C^l} m_j^l = \kappa^l C^l, \quad (1)$$

where  $\Phi(l, i)$  denotes the importance score of the  $i$ -th filter in the  $l$ -th layer,  $\mathbf{m}^l = [m_1^l, m_2^l, \dots, m_{C^l}^l]$  is the mask vector of pruning decision for the  $l$ -th layer.  $m_i^l$  is set to 0 if the  $i$ -th output channel will be pruned or 1 if it will be kept.

### Motivation

From the redundancy perspective, the channel (or filter) set  $\mathcal{T}^l$  should contain those channels whose output features are the most informative ones. In other words, if the output feature distribution of a channel is similar to those of other channels, the information it contains is redundant. In this case, even if this redundant channel is pruned, the next layer can still incorporate similar important information from those channels that survived. Intuitively, channels with high redundancy should be removed, while those less redundant ought to be kept since they contain discriminative information which is difficult to reconstruct from other channels.

Some existing methods like FPGM (He et al. 2019) have explored this idea in image recognition tasks. However, the redundancy (or channel similarity) is calculated based solely on the weights of channels. Despite that these methods achieve promising progress on the image classification task, we find that they fail to achieve satisfactory results on the task of location-sensitive tasks such as semantic segmentation and object detection. This defect can be attributed to the overlook of spatial information. Evaluating the class-related semantic information using the weight of filters has been proven to be effective in image classification models. However, these discrepancies can never capture the variations of spatial distribution between different feature maps, which will lead to poorly located objects for semantic segmentation. The reason for such a performance difference between the pruned classification and segmentation models lies in that classification models typically pool feature maps into vectors, while segmentation models deeply rely on the feature maps with rich spatial information. Therefore, the pruned classification models can still make correct classification predictions based on these spatially misaligned feature maps. Nevertheless, for semantic segmentation, the misaligned feature maps directly divert the attention of succeeding layers from the correct position, which eventually results in poor performance.

To address this issue, we propose a spatial-aware channel redundancy pruning method for semantic segmentation.

In contrast to most existing methods that evaluate the importance of channels, some pioneering works such as CHIP (Sui et al. 2021) and HRank (Lin et al. 2020a) propose to compute the importance score in a data-driven manner: the pre-trained network inference on a set of training data to get real feature maps which are then utilized for the calculation of importance score. Inspired by them, we propose to minimize the feature map information redundancy among the remained channels in the pruned network in a spatial-aware manner.

### Structural Pruning as Maximum Edge Weight Clique Problem

From the perspective of redundancy, each remained channel should produce exclusive discriminate information to maintain low redundancy in the pruned model. In other words, all  $\kappa^l C^l$  channels are selected and kept from all the  $C^l$  ones such that their information redundancy (or similarity) between them is minimized. With this goal, the channel pruning problem can be naturally represented as the *Maximum Edge Weight Clique Problem* (MEWCP) with a cardinality bound. This provides us with a new guideline to design pruning methods in a principled way based on the overall perspective of all the remained channels and graph theory knowledge.

**Definition 1** (*Maximum Edge Weight Clique Problem*). Let  $G = (V, E)$  be a simple, complete, undirected, and edge-weighted graph, where  $V = \{1, 2, \dots, C^l\}$  is the set of channel vertices (nodes) and  $E$  is the set of edges between channel vertices. The clique is defined as a subset of vertices  $C \subseteq V$  which is an induced subgraph of  $G$  that is complete, i.e. for all vertices  $i, j \in C$ , edge  $\{i, j\}$  exist in the corresponding subgraph  $G[C]$ . The weight of the edge  $\{i, j\}$  is denoted as  $a_{ij}$ , which is negatively correlated with the proposed spatial-aware information redundancy metric. The summation of edge weights of clique  $C$  is noted as  $W_C^{edge}$ . The maximum edge weight clique of  $G$  is a clique  $C$  in  $G$  which maximizes  $W_C^{edge}$ . The MEWCP with a cardinality bound constraint aims to find the maximum edge weight clique with no larger than  $b$  vertices, which be formulated as the follows:

$$\max_{\mathbf{m}^l} \sum_{i \in V} \sum_{j \in V \setminus \{i\}} m_i^l m_j^l a_{ij}, \quad s.t. \sum_{i \in V} m_i^l = b. \quad (2)$$

When  $b = \kappa^l C^l$ , taking channels as vertices, and set  $\Phi(l, i) = \sum_{j \in V \setminus \{i\}} m_j^l a_{ij}$ , this problem has the same format with that of channel pruning. The details of edge weight  $a_{ij}$  are introduced in the next subsection.

### Spatial-aware Information Redundancy Metric

To evaluate the information redundancy and keep the spatial information, we utilize the un-pooled feature map for calculation. We have to point out that previous data-driven pruning methods, like HRank and CHIP, only use the pooled vector for importance score calculation, which completely eliminates the spatial information. For semantic segmentation, spatial information in the feature map is vital, but not taken into consideration by existing pruning methods. Thus, we

keep each element in the feature map, preserving the abundant location information for redundancy metric to enable the spatial-aware ability. Here, we propose the redundancy metrics which can be formulated as:

$$r_{ij} = \log 2 - \frac{1}{2} D_{kl}(\mathbf{F}_i^l \parallel \frac{\mathbf{F}_i^l + \mathbf{F}_j^l}{2}) - \frac{1}{2} D_{kl}(\mathbf{F}_j^l \parallel \frac{\mathbf{F}_i^l + \mathbf{F}_j^l}{2}), \quad (3)$$

where  $D_{kl}(\cdot)$  denotes the KL Divergence,  $\mathbf{F}_i^l, \mathbf{F}_j^l \in \mathcal{R}^{C^l \times H \times W}$  denote the un-pooled output feature map of the  $i$ -th and  $j$ -th filter with resolution  $H \times W$ , respectively.

As our optimization goal is to maximize the summation of edge weights, the edge weight should be negatively correlated with the redundancy. Moreover, if the calculation is conducted after pre-training, it would cost a huge memory and time to inference on real training data and collect the feature maps. Thus, we adopt the Exponential Moving Average (EMA) mechanism to accumulate the edge weight since the calculation of  $r_{ij}$  requires high-resolution feature maps of each layer. Finally, the edge weight  $a_{ij}$  can be defined as the follows:

$$a_{ij}^t = a_{ij}^{t-1} \alpha + (1 - r_{ij})(1 - \alpha), \quad (4)$$

where  $\alpha$  is the hyper-parameter of EMA, which is set to 0.99. Hence, with this redundancy metric and edge weight, we can leverage the rich theoretical foundations of MEWCP to derive a spatial-aware pruning method.

### Efficient Heuristic Greedy Pruning

Unfortunately, an practical issue arises when it comes to address this problem in such a channel pruning scenario: MEWCP is apparently NP-hard (Hosseinian et al. 2017), which requires a huge amount of time to get the exact solution. The large number of channels and layers makes it even more time-consuming, since we have to consider all the  $\binom{b}{C^l}$  possibilities for each layer. Although there researchers attempt to reduce the computational complexity, acquiring an exact solution is still too expensive compared to existing channel pruning methods whose complexity is usually only  $\mathcal{O}(C^l + C^l \log(C^l))$ . In addition, the heuristic methods for MEWCP is also not robust enough. Studies shows that the computational complexity of some heuristic methods (Macambira and De Souza 2000) could reach  $\mathcal{O}(n^4)$  or even  $\mathcal{O}(n^5)$  (Hosseinian et al. 2017). Therefore, it is imperative to design a computational efficient algorithm for our MEWCP pruning method.

Although the cost for exactly solving MEWCP is unacceptable for most case, the solution for  $b = C^l - 1$  can be explicitly given (Späth 1985). Putting  $\{k\} := V \setminus C$ , the edge weight of clique  $C$  can be formulated as :

$$\begin{aligned} W_C^{edge} &= \sum_{i \in C} \sum_{j \in C \setminus \{i\}} a_{ij} = \sum_{i \in V} \sum_{j \in V \setminus \{i\}} m_i^l m_j^l a_{ij} \\ &= \sum_{i \in V} \sum_{j \in V \setminus \{i\}} a_{ij} - \sum_{h \in V \setminus \{k\}} a_{hk}. \end{aligned} \quad (5)$$

Thus,  $C$  can be found by only excluding  $k$  out of  $V$  such that the vertices  $k$  has the minimum summation of edge weights:

$$\min_k \sum_{h \in V \setminus \{k\}} a_{hk}. \quad (6)$$

---

**Algorithm 1: Efficient Heuristic Greedy Pruning**

---

**Input** : Weight of the  $l$ -th layer  $\mathbf{W}^l \in \mathcal{R}^{C^l \times C^{l-1}}$  where  $\mathbf{w}_{i,:}^l$  is the weight of  $i$ -th channel, the graph  $G = (V, E)$ , and the channel vertices set  $V = \{1, 2, \dots, C^l\}$

**Output**: The weight  $\mathbf{W}_p^l \in \mathcal{R}^{C_p^l \times C^{l-1}}$  after pruning

- 1 Get the edge weight matrix  $\mathbf{A}^l$
- 2 Set the vertices set  $V_r$  to be removed as empty set
- 3 Set the pruning index  $k = 1$ , counter  $t = 1$   
    // Calculate sum of edge weight:
- 4 **for**  $i \in V$  **do**
- 5      $s_i \leftarrow \sum_{j \in V \setminus \{i\}} a_{ij}$
- 6     **if**  $s_k > s_i$  **then**
- 7          $k \leftarrow i$              // Get the least sum
- // Start iterative pruning:
- 8 **for** each prune iteration  $t_p \in [b]$  **do**
- 9      $V_r \leftarrow V_r + \{k\}$     // Prune the selected  $k$
- 10     $q \leftarrow \text{random}(V - V_r)$
- 11    **for**  $i \in V - V_r$  **do**
- 12        // Update sum for next iteration
- 13         $s_i \leftarrow s_i - a_{ik}$
- 14        **if**  $s_q > s_i$  **then**
- 15             $q \leftarrow i$         // Get the least sum
- 16      $k \leftarrow q$ ;  $t \leftarrow t + 1$
- 17  $\mathbf{W}_p^l \leftarrow \text{Concat}([\mathbf{w}_{i,:}^l, \text{for } i \in V - V_r])$
- 18 **return**  $\mathbf{W}_p^l$

---

Hence, if we prune only one channel out of the  $l$ -th layer, the original optimization problem is equivalent to selecting the node  $k$  with the minimum sum of all its edge weights.

Therefore, we propose a greedy algorithm called efficient heuristic greedy pruning (EHGP) to reduce the complexity of solving MEWCP, as shown in Algorithm 1. We decompose the original problem of pruning  $C^l - b$  channels(nodes) at once into the task of iteratively deleting a single channel for  $C^l - b$  times. First, we calculate the summation of edge weights for each vertices. Then, during each iteration in the EHGP, we prune the single channel vertices with the least sum of edge weights for all layers and update the sum of left edge weights of those still remaining vertices for the next iteration. In this greedy way, our proposed method only needs  $b$  iterations to solve the MEWCP in such a network pruning scenario. Thus, the average complexity of the proposed pruning procedure is reduced to  $\mathcal{O}(n^2)$ . We also compare the runtime of our method and other MEWCP solutions to prune a Deeplabv3-ResNet50 model. The results in Table 6 demonstrate the efficiency of EHGP.

### Pruning Pipeline

As we utilize the feature maps to acquire the information redundancy, a simple way to achieve this is to infer on a set of selected data after the pre-training stage to collect feature maps, which will lead to undesired additional computation

---

**Algorithm 2: Overall Pruning Pipeline**

---

**Input** : An  $L$ -layer model with all the weights  $\Theta = \{\mathbf{W}^1, \dots, \mathbf{W}^L\}$ , and the training dataset  $D$

**Output**: The pruned and finetuned model  $\Theta'_p$

- 1 **for** each pre-training iteration  $t \in [T_{total}^{pre}]$  **do**
- 2     Sample a mini-batch  $(\mathbf{X}, \mathbf{Y})$  from  $D$
- 3     Calculate spatial-aware redundancy based on feature maps             // Equation 3
- 4     Update edge weight matrix  $\mathbf{A}_t^l$  of each layer using EMA             // Equation 4
- 5     Update the unpruned model  $\Theta$
- 6  $\Theta^{prev} \leftarrow \Theta$
- 7 **for** each progressive pruning step  $t_s \in [T_{step}]$  **do**
- 8     Prune  $\Theta^{prev}$  by EHGP in Algorithm 1 and get pruned model  $\Theta_p^{t_s} = \{\mathbf{W}_p^1, \dots, \mathbf{W}_p^L\}^{t_s}$
- 9     **for** each fine-tuning iteration  $t \in [T_{total}^{fine}]$  **do**
- 10        Sample a mini-batch  $(\mathbf{X}, \mathbf{Y})$  from  $D$
- 11        Update the weight of pruned model  $\Theta_p^{t_s}$
- 12     Set the pruned and finetuned model  $\Theta_f^{t_s}$  as the initial model for the next step:  $\Theta^{prev} = \Theta_f^{t_s}$
- 13 **return** the final pruned and finetuned model  $\Theta_f^{T_{step}}$ .

---

costs. Some recent works (Lin et al. 2020a; Sui et al. 2021) adopt this *pretrain-collect-prune-finetune* pipeline. Accordingly, they suffer from extremely large time costs (more than an hour on ImageNet (Deng et al. 2009)) to perform the *collect&prune*. Inspired by previous works (Ioffe and Szegedy 2015) we adopt the Exponential Moving Average (EMA) mechanism to accumulate the edge weight matrix  $\mathbf{A}^l$  during training, as shown in Equation 4. In order to mitigate the large performance drop after pruning, we progressively remove channels in  $T_{step}$  iterations. We prune the model and finetune it for a few iterations to restore the performance, which makes the pruning less destructive. This can also reduce the duration of accumulating edge weight matrix and lead to an accurate reflection of the feature redundancy between channels. The detailed overall pruning process is shown in the Algorithm 2.

### Experiments and Analysis

Although this paper mainly focuses on semantic segmentation, to comprehensively evaluate the efficacy and generality of our method, we also conduct experiments on a variety of computer vision tasks, including image classification and object detection. The experiments run on the Pytorch framework using NVIDIA RTX 3090 GPUs. More implementation details, ablation studies, and visualization results are presented in the Appendix.

### Datasets

In this work, we conduct experiments on three tasks: semantic segmentation, image classification, and object detec-

Method	mIoU(%)	Params(%↓)	FLOPs(%↓)
Unpruned	81.6	41.27M	1418.29G
Random	78.7	16.60M (59.78%↓)	558.72G (60.61%↓)
NS	79.9	15.57M (62.27%↓)	575.20G (59.44%↓)
Taylor	80.3	14.99M (63.68%↓)	564.88G (60.17%↓)
DepGraph	80.0	16.84M (59.20%↓)	561.92G (60.38%↓)
FPGM-60%	80.2	14.68M (64.43%↓)	550.75G (61.17%↓)
DCFP-60%	80.9	14.78M (64.19%↓)	553.74G (60.96%↓)
<b>Ours-60%</b>	<b>81.3</b>	<b>14.52M (64.82%↓)</b>	<b>549.84G (61.23%↓)</b>
FPGM-70%	79.3	10.51M (74.82%↓)	411.26G (71.00%↓)
DCFP-70%	79.8	10.60M (74.32%↓)	418.08G (70.52%↓)
<b>Ours-70%</b>	<b>80.9</b>	<b>10.45M (74.44%↓)</b>	<b>398.56G (71.90%↓)</b>
FPGM-80%	77.9	6.27M (84.81%↓)	274.18G (80.67%↓)
DCFP-80%	78.8	6.65M (83.89%↓)	280.96G (80.19%↓)
<b>Ours-80%</b>	<b>79.4</b>	<b>5.88M (85.75%↓)</b>	<b>260.99G (81.60%↓)</b>

Table 1: Comparison with other pruning methods on Cityscapes validation set using Deeplabv3-ResNet50.

Method	Deeplabv3+		PSPNet	
	mIoU(%)	FLOPs(%↓)	mIoU(%)	FLOPs
Unpruned	81.9	1608.99G	81.3	1476.54G
Random	76.4	666.40G	76.8	582.08G
DCFP-60%	80.6	652.08G	79.1	595.04G
<b>Ours-60%</b>	<b>81.6</b>	<b>614.81G</b>	<b>81.1</b>	<b>573.96G</b>
DCFP-70%	79.5	502.32G	78.1	444.16G
<b>Ours-70%</b>	<b>81.2</b>	<b>460.02G</b>	<b>80.2</b>	<b>412.52G</b>
DCFP-80%	78.4	330.96G	76.0	314.40G
<b>Ours-80%</b>	<b>80.3</b>	<b>306.01G</b>	<b>79.6</b>	<b>286.93G</b>

Table 2: Ablation studies on Deeplabv3+ and PSPNet. All models are trained on Cityscapes.

tion. Following DCFP, aside from the popular but rather easy Cityscapes (Cordts et al. 2016) dataset (19 categories), we also evaluate our method on the challenging ADE20k (Zhou et al. 2017) (150 categories) and COCO stuff-10K (Caesar, Uijlings, and Ferrari 2018) (171 categories). For image classification and object detection, we compare the proposed method with other state-of-the-art pruning methods on ImageNet (Deng et al. 2009) and COCO2017 dataset (Lin et al. 2014), respectively.

## Semantic Segmentation

**Comparison on Cityscapes** To compare our method with other popular pruning algorithms including NS (Liu et al. 2017), Taylor (Molchanov et al. 2019), Depgraph (Fang et al. 2023), FPGM (He et al. 2019), and DCFP (Wang et al. 2024), we reproduce these methods and prune the DeeplabV3-resnet50 model on Cityscapes. The results are shown in Table 1. Following DCFP, all our experiments adopt the GSRL to improve performance. Our method attains a significant performance advantage over the other pruning methods. Moreover, we also include recent real-time segmentation methods and more pruning methods for segmentation like CAP (He et al. 2021) and MTP (Chen, Zhang, and Wang 2022) in Figure 1 to compare the number of parameters and mIoU performance. According to the results, SIRFP achieves the best parameter-mIoU trade-off

Method	ADE20K		COCO stuff-10k	
	mIoU(%)	FLOPs	mIoU(%)	FLOPs
Unpruned	45.3	177.42G	38.1	177.45G
Random	40.2	71.29G	28.3	70.59G
NS	40.6	72.33G	31.1	71.61G
Taylor	40.8	73.83G	32.0	70.89G
DepGraph	41.0	70.31G	32.0	70.89G
FPGM	39.8	72.44G	30.3	72.44G
DCFP-60%	44.8	71.10G	34.3	70.45G
<b>Ours-60%</b>	<b>44.9</b>	<b>67.82G</b>	<b>34.8</b>	<b>69.65G</b>
DCFP-70%	44.4	55.93G	33.1	51.92G
<b>Ours-70%</b>	<b>44.6</b>	<b>51.40G</b>	<b>33.9</b>	<b>50.74G</b>
DCFP-80%	41.4	35.96G	30.3	35.83G
<b>Ours-80%</b>	<b>41.8</b>	<b>32.39G</b>	<b>30.7</b>	<b>33.25G</b>

Table 3: Results comparison with other pruning methods of Deeplabv3-ResNet50 on ADE20K and COCO stuff-10K.

Method	FLOPs Reduction	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
Baseline	0%	25.2	42.7	25.8	7.3	27.1	40.8
DMCP	50%	24.1	41.2	24.7	6.7	25.6	39.2
FPGM*	50%	25.2	42.4	25.9	7.0	27.2	41.3
CHEX-1	50%	25.9	43.0	26.8	7.8	27.8	41.7
<b>Ours</b>	50%	<b>26.6</b>	<b>44.2</b>	<b>27.3</b>	<b>8.1</b>	<b>28.8</b>	<b>43.4</b>
CHEX-2	75%	24.3	41.0	24.9	7.1	25.6	40.1
<b>Ours</b>	75%	<b>24.6</b>	<b>41.8</b>	<b>25.1</b>	<b>7.3</b>	<b>26.0</b>	<b>41.2</b>

Table 4: Object detection pruning results on COCO2017 validation set using SSD-ResNet50. \* denotes our reproduction.

among all the other pruning methods and real-time methods. **Comparison on ADE20K and COCO stuff-10k** The comparison with existing pruning methods on ADE20K and COCO stuff-10K are reported in Table 3, showing consistent effectiveness on these challenging datasets. Specifically, our method exceeds the FPGM by 1.1%, 5.0%, and 4.5% mIoU under 60% FLOPs reduction on Cityscapes, ADE20K, and COCO stuff-10k, respectively. We attribute this significant performance advantage to the proposed spatial-aware redundancy metric, which enables our method to capture the overlap of spatial distribution between different features.

**Comparison with DCFP** As DCFP (Wang et al. 2024) also calculates the importance score using the EMA mechanism, the performance advantage of our method over DCFP indicates that the adoption of the EMA mechanism is obviously not the main reason for performance advantage. We believe that this corroborates our argument that the spatial-aware metric is more suitable for semantic segmentation than traditional spatial-agnostic metrics, like the grad-based importance score of DCFP.

## Object Detection

Object detection is also a location-sensitive computer vision task. We also conduct experiments on the classic SSD-ResNet50 and compare the results with DMCP (Guo et al. 2020), FPGM (He et al. 2019), CHEX (Hou et al. 2022) to

Method	Top-1	Top-5	FLOPs
SSS (Huang and Wang 2018)	74.18	91.91	2.8G
GReg-2 (Wang et al. 2020)	75.36	-	2.8G
AOFP-C1 (Ding et al. 2019)	75.63	92.69	2.6G
ThiNet (Luo, Wu, and Lin 2017)	72.04	90.67	2.4G
SFP (He et al. 2018)	74.61	92.87	2.4G
FPGM (He et al. 2019)	75.50	92.63	2.4G
BNFI (Oh et al. 2022)	75.47	-	2.3G
TAS (Dong and Yang 2019)	76.20	-	2.3G
HRank (Lin et al. 2020a)	74.98	92.33	2.3G
Taylor (Molchanov et al. 2019)	74.50	-	2.3G
CCP (Peng et al. 2019)	75.50	92.62	2.1G
GFP (Liu et al. 2021a)	76.42	-	2.0G
DepGraph (Fang et al. 2023)	75.28	-	2.0G
DSA (Ning et al. 2020)	74.69	92.06	2.0G
CafeNet (Su et al. 2021)	76.90	93.30	2.0G
<b>Ours</b>	<b>77.35</b>	<b>93.60</b>	<b>2.0G</b>
BNFI (Oh et al. 2022)	75.02	-	1.9G
FPGM (He et al. 2019)	74.83	92.32	1.9G
Hinge (Li et al. 2020)	74.70	-	1.9G
DCP (Zhuang et al. 2018)	74.95	<b>93.54</b>	1.8G
<b>Ours</b>	<b>75.14</b>	93.12	<b>1.7G</b>
HRank (Lin et al. 2020a)	71.98	91.01	1.6G
DMCP (Guo et al. 2020)	74.10	-	1.1G
CHIP (Sui et al. 2021)	73.30	91.48	1.0G
<b>Ours</b>	<b>74.67</b>	<b>92.19</b>	<b>1.0G</b>

Table 5: Comparison with existing pruning methods on ImageNet validation set using ResNet-50.

evaluate the effectiveness and generalization of our method. Experimental results conducted on the COCO2017 dataset are reported in Table 4. Following CHEX, we only prune the backbone of SSD. Compared with the data-driven method CHEX which relies on the rank of the feature map, our method achieves a 0.7% AP improvement over CHEX under 50% FLOPs reduction ratio. The performance advantages over our reproduced FPGM and CHEX illustrate that our spatial-aware method is also more powerful than the other spatial-agnostic pruning methods on object detection task which also requires excellent localization ability.

## Image Classification

Although we focus on location-sensitive tasks, the performance on image classification also indicates the ability to preserve spatial-agnostic semantic information. Hence, we conduct experiments on the ImageNet dataset and the results are reported in Table 5. Following the adaptive pruning methods, we prune the model during training progressively. SIRFP obtains a significant performance advantage over existing pruning methods. As CHIP (Sui et al. 2021) and HRank (Lin et al. 2020a) also utilize features to calculate importance score, we contend that the spatial-aware redundancy metric, instead of the data-driven manner, is the main cause of our performance advantage.

## Ablation Studies

**Generalization on different models.** To evaluate the effectiveness and generalization ability of our pruning method,

Method	Type	Runtime(sec.)	mIoU(%)
(Pullan 2008)	Heuristic	3731	81.3
(Macambira et al. 2000)	Heuristic	4933	81.1
(Hosseinian et al. 2020)	Exact	>10800	81.4
<b>EHGP(Ours)</b>	Heuristic	285	81.3

Table 6: Runtime comparison between our proposed EHGP and different methods for the MEWCP. Experiments are conducted on Deeplabv3-ResNet50.

Method	FLOPs Reduction	FPS(img/s)	mIoU(%)
Unpruned	0%	36.7	79.3
DCFP	60%	61.7	80.7
<b>Ours</b>	60%	<b>62.4</b>	<b>81.3</b>
DCFP	70%	69.9	79.8
<b>Ours</b>	70%	<b>73.9</b>	<b>80.9</b>
DCFP	80%	81.1	78.8
<b>Ours</b>	80%	<b>90.5</b>	<b>79.4</b>

Table 7: Performance and inference speed comparison with DCFP on Cityscapes dataset using Deeplabv3-ResNet50.

we conduct experiments on various segmentation models. Aside from Deeplabv3 (Chen et al. 2017a), we also conduct experiments on the popular and classic PSPNet (Zhao et al. 2017) and Deeplabv3+ (Chen et al. 2018), as reported in the Table 2. The experimental results demonstrate the decent generalization ability of SIRFP on various models.

**Ablation on MEWCP solver.** In order to show the efficiency of our greedy pruning algorithm, we compare the runtime with other exact and heuristic methods for MEWCP in the pruning scenario. The results in Table 6 show that our greedy algorithm exhibits outstanding efficiency compared with other solutions. With only  $\mathcal{O}(n^2)$  complexity as we discussed above, the speed of the proposed EHGP significantly exceeds other methods for MEWCP.

**Time efficiency comparison** We compare the inference time efficiency of DCFP and our SIRFP in Table 7. All our models employ TensorRT, the commonly used SDK for real-time neural network inference. Compared with DCFP, our method performs more efficiently with higher mIoU results and faster inference speed, especially on larger FLOPs reduction ratios.

## Limitation

Although SIRFP exhibits excellent performance and generalization ability, it relies on the EMA to accumulate edge weight matrix during training. For a given well-trained network, the SIRFP has to inference on training samples to acquire feature maps, causing extra computational cost. Moreover, the time to calculate redundancy may be huge for networks whose feature map resolution is extremely large. Thus, the efficiency of redundancy calculation can be further improved.

## Acknowledgments

This work was supported by the Hubei Provincial Natural Science Foundation of China No.2022CFA055, the National Natural Science Foundation of China No.62176097, and Interdisciplinary Research Program of HUST No.2024JCYJ034.

## References

- Caesar, H.; Uijlings, J.; and Ferrari, V. 2018. COCO-Stuff: Thing and Stuff Classes in Context. In *CVPR*, 1209–1218.
- Carreira-Perpinán, M. A.; and Idelbayev, Y. 2018. “learning-compression” algorithms for neural net pruning. In *CVPR*, 8532–8541.
- Chen, L.-C.; Papandreou, G.; Kokkinos, I.; Murphy, K.; and Yuille, A. L. 2014. Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062*.
- Chen, L.-C.; Papandreou, G.; Kokkinos, I.; Murphy, K.; and Yuille, A. L. 2017a. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *TPAMI*, 40(4): 834–848.
- Chen, L.-C.; Papandreou, G.; Schroff, F.; and Adam, H. 2017b. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*.
- Chen, L.-C.; Zhu, Y.; Papandreou, G.; Schroff, F.; and Adam, H. 2018. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 801–818.
- Chen, X.; Zhang, Y.; and Wang, Y. 2022. MTP: multi-task pruning for efficient semantic segmentation networks. In *ICME*, 1–6.
- Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; and Schiele, B. 2016. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 3213–3223.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*, 248–255.
- Ding, H.; Jiang, X.; Shuai, B.; Liu, A. Q.; and Wang, G. 2018. Context contrasted feature and gated multi-scale aggregation for scene segmentation. In *CVPR*, 2393–2402.
- Ding, X.; Ding, G.; Guo, Y.; Han, J.; and Yan, C. 2019. Approximated oracle filter pruning for destructive cnn width optimization. In *ICML*, 1607–1616.
- Dong, X.; and Yang, Y. 2019. Network pruning via transformable architecture search. *NeurIPS*, 32.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Esser, S. K.; McKinstry, J. L.; Bablani, D.; Appuswamy, R.; and Modha, D. S. 2019. Learned step size quantization. *arXiv preprint arXiv:1902.08153*.
- Fan, M.; Lai, S.; Huang, J.; Wei, X.; Chai, Z.; Luo, J.; and Wei, X. 2021. Rethinking bisenet for real-time semantic segmentation. In *CVPR*, 9716–9725.
- Fang, G.; Ma, X.; Song, M.; Mi, M. B.; and Wang, X. 2023. Depgraph: Towards any structural pruning. In *CVPR*, 16091–16101.
- Guo, S.; Wang, Y.; Li, Q.; and Yan, J. 2020. Dmcp: Differentiable markov channel pruning for neural networks. In *CVPR*, 1539–1547.
- Han, S.; Mao, H.; and Dally, W. J. 2015. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *arXiv preprint arXiv:1510.00149*.
- Han, S.; Pool, J.; Tran, J.; and Dally, W. 2015. Learning both weights and connections for efficient neural network. *NeurIPS*, 28.
- He, W.; Wu, M.; Liang, M.; and Lam, S.-K. 2021. Cap: Context-aware pruning for semantic segmentation. In *WACV*, 960–969.
- He, Y.; Kang, G.; Dong, X.; Fu, Y.; and Yang, Y. 2018. Soft Filter Pruning for Accelerating Deep Convolutional Neural Networks. In *IJCAI*, 2234–2240.
- He, Y.; Liu, P.; Wang, Z.; Hu, Z.; and Yang, Y. 2019. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *CVPR*, 4340–4349.
- Hosseinian, S.; Fontes, D. B.; and Butenko, S. 2020. A Lagrangian bound on the clique number and an exact algorithm for the maximum edge weight clique problem. *INFORMS Journal on Computing*, 32(3): 747–762.
- Hosseinian, S.; Fontes, D. B.; Butenko, S.; Nardelli, M. B.; Fornari, M.; and Curtarolo, S. 2017. The maximum edge weight clique problem: formulations and solution approaches. *Optimization Methods and Applications: In Honor of Ivan V. Sergienko’s 80th Birthday*, 217–237.
- Hou, Z.; Qin, M.; Sun, F.; Ma, X.; Yuan, K.; Xu, Y.; Chen, Y.-K.; Jin, R.; Xie, Y.; and Kung, S.-Y. 2022. Chex: Channel exploration for cnn model compression. In *CVPR*, 12287–12298.
- Huang, Z.; and Wang, N. 2018. Data-driven sparse structure selection for deep neural networks. In *ECCV*, 304–320.
- Ioffe, S.; and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 448–456.
- Li, H.; Kadav, A.; Durdanovic, I.; Samet, H.; and Graf, H. 2016. Pruning Filters for Efficient ConvNets. *arXiv preprint arXiv:1608.08710*.
- Li, Y.; Gu, S.; Mayer, C.; Gool, L. V.; and Timofte, R. 2020. Group sparsity: The hinge between filter pruning and decomposition for network compression. In *CVPR*, 8018–8027.
- Lin, M.; Ji, R.; Wang, Y.; Zhang, Y.; Zhang, B.; Tian, Y.; and Shao, L. 2020a. Hrank: Filter pruning using high-rank feature map. In *CVPR*, 1529–1538.
- Lin, T.; Stich, S.; Barba, L.; Dmitriev, D.; and Jaggi, M. 2020b. Dynamic Model Pruning with Feedback. *ICLR*.

- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *ECCV*, 740–755.
- Liu, L.; Zhang, S.; Kuang, Z.; Zhou, A.; Xue, J.-H.; Wang, X.; Chen, Y.; Yang, W.; Liao, Q.; and Zhang, W. 2021a. Group fisher pruning for practical network compression. In *ICML*, 7021–7032.
- Liu, Z.; Li, J.; Shen, Z.; Huang, G.; Yan, S.; and Zhang, C. 2017. Learning efficient convolutional networks through network slimming. In *CVPR*, 2736–2744.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021b. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 10012–10022.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *CVPR*, 3431–3440.
- Luo, J.-H.; Wu, J.; and Lin, W. 2017. Thinet: A filter level pruning method for deep neural network compression. In *CVPR*, 5058–5066.
- Macambira, E. M.; and De Souza, C. C. 2000. The edge-weighted clique problem: valid inequalities, facets and polyhedral computations. *European Journal of Operational Research*, 123(2): 346–371.
- Molchanov, P.; Mallya, A.; Tyree, S.; Frosio, I.; and Kautz, J. 2019. Importance estimation for neural network pruning. In *CVPR*, 11264–11272.
- Ning, X.; Zhao, T.; Li, W.; Lei, P.; Wang, Y.; and Yang, H. 2020. DSA: More Efficient Budgeted Pruning via Differentiable Sparsity Allocation. In *ECCV*, 592–607.
- Oh, J.; Kim, H.; Baik, S.; Hong, C.; and Lee, K. M. 2022. Batch normalization tells you which filter is important. In *WACV*, 2645–2654.
- Pan, H.; Hong, Y.; Sun, W.; and Jia, Y. 2022. Deep dual-resolution networks for real-time and accurate semantic segmentation of traffic scenes. *T-ITS*, 24(3): 3448–3460.
- Pang, Y.; Li, Y.; Shen, J.; and Shao, L. 2019. Towards bridging semantic gap to improve semantic segmentation. In *CVPR*, 4230–4239.
- Peng, H.; Wu, J.; Chen, S.; and Huang, J. 2019. Collaborative channel pruning for deep networks. In *ICML*, 5113–5122.
- Pullan, W. 2008. Approximating the maximum vertex/edge weighted clique using local search. *Journal of Heuristics*, 14: 117–134.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, 234–241.
- Späth, H. 1985. Heuristically determining cliques of given cardinality and with minimal cost within weighted complete graphs. *Zeitschrift für Operations Research*, 29: 125–131.
- Su, X.; You, S.; Huang, T.; Wang, F.; Qian, C.; Zhang, C.; and Xu, C. 2021. Locally Free Weight Sharing for Network Width Search. *ICLR*.
- Sui, Y.; Yin, M.; Xie, Y.; Phan, H.; Aliari Zonouz, S.; and Yuan, B. 2021. Chip: Channel independence-based pruning for compact neural networks. *NeurIPS*, 34: 24604–24616.
- Wan, Q.; Huang, Z.; Lu, J.; Yu, G.; and Zhang, L. 2023. Seaformer: Squeeze-enhanced axial transformer for mobile semantic segmentation. *arXiv preprint arXiv:2301.13156*.
- Wang, H.; Qin, C.; Zhang, Y.; and Fu, Y. 2020. Neural pruning via growing regularization. *arXiv preprint arXiv:2012.09243*.
- Wang, J.; Gou, C.; Wu, Q.; Feng, H.; Han, J.; Ding, E.; and Wang, J. 2022. RTFormer: Efficient design for real-time semantic segmentation with transformer. *NeurIPS*, 35: 7423–7436.
- Wang, Z.; Xie, H.; Wang, Y.; Xu, H.; and Jin, G. 2024. DCFP: Distribution Calibrated Filter Pruning for Lightweight and Accurate Long-tail Semantic Segmentation. *TCSVT*, 34(7): 6063–6076.
- Xu, Y.; Wang, Y.; Zhou, A.; Lin, W.; and Xiong, H. 2018. Deep neural network compression with single and multiple level quantization. In *AAAI*, volume 32.
- Yang, Z.; Wang, Y.; Liu, C.; Chen, H.; Xu, C.; Shi, B.; Xu, C.; and Xu, C. 2019. Legonet: Efficient convolutional neural networks with lego filters. In *ICML*, 7005–7014.
- Yu, C.; Gao, C.; Wang, J.; Yu, G.; Shen, C.; and Sang, N. 2021. Bisenet v2: Bilateral network with guided aggregation for real-time semantic segmentation. *IJCV*, 129: 3051–3068.
- Yu, C.; Wang, J.; Peng, C.; Gao, C.; Yu, G.; and Sang, N. 2018. Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *ECCV*, 325–341.
- Yu, X.; Liu, T.; Wang, X.; and Tao, D. 2017. On compressing deep models by low rank and sparse decomposition. In *CVPR*, 7370–7379.
- Zhao, H.; Shi, J.; Qi, X.; Wang, X.; and Jia, J. 2017. Pyramid scene parsing network. In *CVPR*, 2881–2890.
- Zhou, B.; Zhao, H.; Puig, X.; Fidler, S.; Barriuso, A.; and Torralba, A. 2017. Scene parsing through ade20k dataset. In *CVPR*, 633–641.
- Zhuang, Z.; Tan, M.; Zhuang, B.; Liu, J.; Guo, Y.; Wu, Q.; Huang, J.; and Zhu, J. 2018. Discrimination-aware channel pruning for deep neural networks. *NeurIPS*, 31.