arXiv:2101.06686v1 [cs.CV] 17 Jan 2021

KCP: Kernel Cluster Pruning for Dense Labeling Neural Networks

Po-Hsiang Yu
National Taiwan University
r08943024@ntu.edu.tw

Sih-Sian Wu
National Taiwan University
benwu@video.ee.ntu.edu.tw

Liang-Gee Chen
National Taiwan University
lgchen@ntu.edu.tw

Abstract

Pruning has become a promising technique used to compress and accelerate neural networks. Existing methods are mainly evaluated on spare labeling applications. However, dense labeling applications are those closer to real world problems that require real-time processing on resource-constrained mobile devices. Pruning for dense labeling applications is still a largely unexplored field. The prevailing filter channel pruning method removes the entire filter channel. Accordingly, the interaction between each kernel in one filter channel is ignored.

In this study, we proposed kernel cluster pruning (KCP) to prune dense labeling networks. We developed a clustering technique to identify the least representational kernels in each layer. By iteratively removing those kernels, the parameter that can better represent the entire network is preserved; thus, we achieve better accuracy with a decent model size and computation reduction. When evaluated on stereo matching and semantic segmentation neural networks, our method can reduce more than 70% of FLOPs with less than 1% of accuracy drop. Moreover, for ResNet-50 on ILSVRC-2012, our KCP can reduce more than 50% of FLOPs reduction with 0.13% Top-1 accuracy gain. Therefore, KCP achieves state-of-the-art pruning results.

1. Introduction

Deep learning based algorithms have been evolving to achieve significant results in solving complex and various computer vision applications, such as image classification [14, 20, 35], stereo matching [41, 2], and semantic segmentation [27, 3, 37]. However, real world applications often require the deployment of these algorithms on resource-limited mobile/edge devices. Owing to over-parameterization and high computational cost, it is difficult to run those cumbersome models on small devices [40], thereby hindering the fruits of those state-of-the-art (SOTA) intelligent algorithms from benefiting our lives. Neural network pruning is a technique to strike the balance between accuracy (performance) and cost (efficiency). Filter level

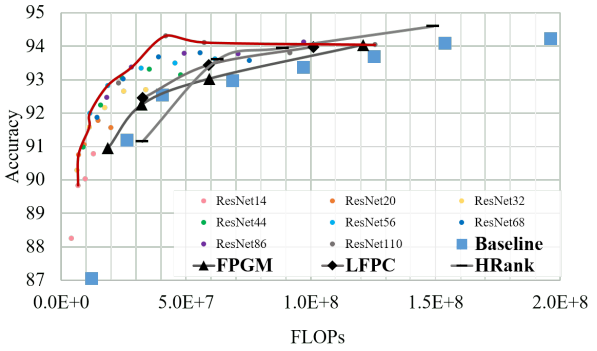

Figure 1. KCP benchmarked on CIFAR-10. Points that is located near upper left are preferred because they achieved better accuracy-FLOPs trade-off. Baseline points are the original ResNet. Lines are the envelopes of points from different depth ResNet. Our method achieves better efficiency than previous state-of-the-arts, namely, FPGM [17], LFPC [15], and HRank [24]. Model pruned by KCP are closer to the upper left corner, and our envelope is decently above that of all other works.

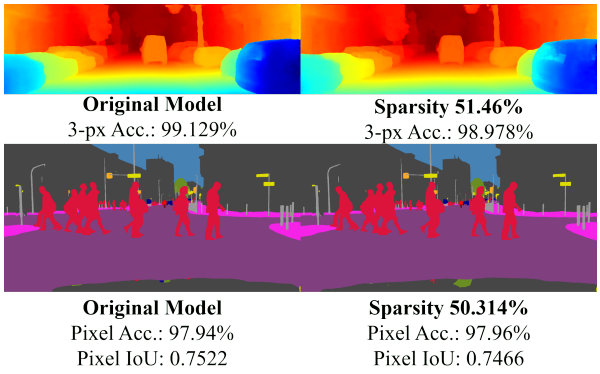

Figure 2. Results of KCP on pruning dense labeling networks. The upper row show the results of depth estimation; the predicted depth are samples from KITTI2015 [30, 31] generated by the original PSMNet [2] and the KCP processed 50% sparsity model. The lower row shows the result of semantic segmentation; the predicted class label are samples from Cityscapes [5] generated by original HRNet [37] and KCP processed 50% sparsity model. KCP maintains the model performance both quantitatively and qualitatively.

pruning has been proven to be an effective [23, 29, 17, 24, 10, 15] and favorable method because it could result in a more structural model.

Despite the numerous filter pruning algorithms that have been proposed, previous works are mostly evaluated on a sparse labeling classification task. In a practical scenario, dense labeling problems often require real-time processing in real world problems. Hence, the compression of a dense labeling neural network is crucial and demanding, yet this field is still quite unexplored. Moreover, filter channel pruning removes the entire output channel, which can sometimes lead to sub-optimal results in terms of accuracy preserving. A single filter channel often contains several kernels; thus, the information density is still high. If one removes the entire channel, the interactions between each kernels are neglected. In such a situation, each kernel in one filter may contribute differently to the final prediction. Pruning at the level of the channel is equivalent to aggregating the influence of an individual kernel, which dilutes the uniqueness of each kernel.

We propose kernel cluster pruning (KCP) to address the issues and fill the gap in dense labeling pruning. The main concept of KCP is to identify the least representational kernels in each layer and subsequently remove them. The smallest target to be pruned in KCP is the kernel instead of the filter. First, the cluster center of kernels in each layer are calculated. Thereafter, the kernels that are closest to the cluster center are considered to carry less information and are removed. Because every pixel in the prediction is assigned a label, dense labeling networks are more vulnerable to change in parameter than sparse labeling networks. Kernel pruning enables us to compress dense labeling networks more delicately while still maintaining the network structure after pruning. Experiments showed that we not only successfully compressed the dense labeling network but also achieved SOTA results on CIFAR-10 [19] and ImageNet (ILSVRC-2012) [34].

Our main contributions are summarized as follows:

- A novel kernel pruning method, KCP, is proposed. This method can identify the kernels with least representativeness in each layer then prune those kernels without breaking the original network structure.
- Thorough evaluations of dense labeling neural network pruning are presented. The results show that KCP can effectively prune those networks and retain the accuracy. To the best of our knowledge, this is the first work that has investigated structured filter pruning on dense labeling applications.
- The experiment on benchmark datasets demonstrates the effectiveness of KCP. KCP achieves new SOTA results for ResNet on CIFAR-10 (c.f. Fig 1, Table 1) and ImageNet (c.f. Table 2).

2. Related Works

The prior art on network pruning can be approximately divided into several approaches, depending on the pruning objectives and problem formulation. The early pruning methods are individual *weight pruning*. The idea can first be observed in optimal brain damage [6] and optimal brain surgeon [13]. More recently, Han *et al*. [12] proposed a series of works removing connections with weights below a certain threshold and later incorporated the idea into the deep compression pipeline [11]. In [26] (Liu *et al*.), the results indicated that the value of automatic structured pruning algorithms sometimes lie in identifying efficient structures and performing an implicit architecture search, rather than selecting "important" weights. Different from rethinking [26] ,the "lottery ticket hypothesis", proposed by Frankle *et al*. [8] is another interesting work about the reason for pruning. They found that a standard pruning technique naturally uncovers sub-networks whose initializations made them capable of being trained effectively.

Meanwhile, *structured* pruning methods prune the entire channel or filter away. Channel pruning has become one of the most popular techniques for pruning owing to its fine-grained nature and ease of fitting into the conventional deep learning framework. Some heuristic approaches use a magnitude-based criterion to determine whether the filter is important. Li *et al*. [23] calculated the $L1$-norm of filters as an important score and directly pruned the corresponding unimportant filter. Luo *et al*. [29] selected unimportant filters to prune based on the statistics information computed from its subsequent layer. He *et al*. [16] proposed to allow pruned-away filters to recover during fine-tuning while still using $L1$-norm criterion to select the filter. Molchanov *et al*. [32] further proposed a method that measures squared loss change induced by removing a specific neuron, which can be applied to any layer. He *et al*. [17] calculated the geometric median of the filters in the layer and removed the filter closest to it. Lin *et al*. [24] discovered that the average rank of multiple feature maps generated by a single filter is always the same, regardless of the number of image batches convolution neural networks (CNNs) receive. They explored the high rank of feature maps and removed the low-rank feature maps to prune a network. He *et al*. [15] developed a differentiable pruning criteria sampler to select different pruning criteria for different layers. Group sparsity imposed on sparsity learning was also extensively used to learn the sparse structure while training a network. Some used $L1$-regularization (Liu *et al*.) [25] term in the training loss, whereas others imposed group LASSO (Wen *et al*.) [38]. Ye *et al*. [39] similarly enforced sparsity constraint on channels during training, and the magnitude would be later used to prune the constant value channels. Lemaire *et al*. [22] introduced a budgeted regularized pruning framework that set activation vol-

ume as a budget-aware regularization term in the objective function. Luo *et al*. [28] focused on compressing residual blocks via a Kullback-Leibler (KL) divergence based criterion along with combined data augmentation and knowledge distillation to prune with limited-data. Guo *et al*. [10] modeled channel pruning as a Markov process. Chin *et al*. [4] proposed the learning of a global ranking of the filter across different layers of the CNN. Gao *et al*. [9] utilized a discrete optimization method to optimize a channel-wise differentiable discrete gate with resource constraint to obtain a compact CNN with string discriminative power.

Discussion. As above mentioned, pruning technique is extensive; however, those works are almost exclusively evaluated on sparse labeling classification tasks, namely, ResNet [14] on CIFAR-10 [19] and ImageNet [34]. To the best of our knowledge, pruning for dense labeling networks is still a largely unexplored field. Yu *et al*. [40] proposed a prune and quantize pipeline for stereo matching networks, yet their objective was different from ours. They sought extremely high sparsity for direct hardware implementation; thus, they applied **norm-based individual weight pruning** and quantized the remaining weights. The resulting model of [40] was **unstructured**, which required dedicated hardware to achieve the claimed ideal acceleration. Our proposed method prunes the model **structurally**, which implies that no special hardware is needed for ideal compression and acceleration.

3. Proposed Method

In this section we will present KCP, which can remove convolutional kernels whose representativeness is lower in dense labeling networks while maintaining a regular network structure.

3.1. Preliminaries (Notation Definitions)

We introduce notation and symbols that will be used throughout this part. We suppose that a neural network has L layers. C^j_{in} and C^j_{out} are the number of input and output channels for the jth convolution layer in the network, respectively. The dimension of the jth convolution layer is $\mathbb{R}^{C^j_{in}\times C^j_{out}\times K^j\times K^j}$, where K is the kernel size of the filter. Thus, in the jth convolution layer, there are C^j_{out} filters ($\mathcal{F}$) and $C^j_{in}\times C^j_{out}$ kernels ($\mathcal{K}$). $\mathcal{K}^j_{m,n}$ is the mth kernel in the nth filter of jth layer. Therefore, we can use kernels to represent the jth layer ($\mathbf{W}^{(\mathrm{j})}$) of the network as $\{\mathcal{K}^j_{m,n}$, $1\le m\le C^j_{in}$, and $1\le n\le C^j_{out}$ while $m,n\in\mathbb{N}\}$.

3.2. Kernel Cluster Pruning

Recent filter pruning methods prune at the level of channels and the target of the pruning unit is one filter. While filter channel pruning has garnered attention in the recent research field owing to its resultant structured sparsity, removing the entire filter output channel sometimes makes it hard to identify the optimal pruning target. There are still several kernel weights in one filter. If one removes the entire filter, the interactions between each kernel in that filter are neglected and this would often result in worse accuracy after pruning. This is empirically proven in Sec. 4.2.

Neural networks designed for solving dense labeling applications often utilize different resolutions of input to obtain better prediction. Such practices leads to the need for a large receptive field in convolutional layers and requires well-designed up/downsampling layers. The performance of these networks is evaluated using pixel-wise metrics, which is prone to changes in parameters compared with typical accuracy for classification tasks.

To solve the problem of filter pruning and dense labeling network compression while maintaining the regular network structure, we proposed KCP, a novel method that prunes **kernels** as the target instead of the filter. Our method is inspired by the concept of K-means clustering. Given a set of n points as follows: $x_i\in\mathbb{R}^d, i=1,2,3,...n$ and k clusters as follows: $\{S_1,S_2,...S_k \quad k\le n\}$. K-means aims to select the centroid that minimizes inertia between the data in the clusters and the cluster centers as follows:

$$\arg\min_{\mu\in\mathbb{R}^d}\sum_{c=1}^{k}\sum_{i=1}^{n_c}\|x_i-\mu_c\|^2\Big|_{x_i\in S_c} \tag{1}$$

μ_c denotes the cluster center and $\|x-\mu\|$ the Euclidean distance. As the inertia measures how internally coherent the clusters are, we use the cluster center to extract the information between each kernel within each layer. In the jth layer, set the number of clusters $k=1$; the cluster center of that layer $\mu^j\in\mathbb{R}^{K^j\times K^j}$ is as follows:

$$\mu^j=\frac{1}{C^j_{in}C^j_{out}}\sum_{m=1}^{C^j_{in}}\sum_{n=1}^{C^j_{out}}\mathcal{K}^j_{m,n} \tag{2}$$

The cluster center anchors the aggregate kernel representativeness of that layer in the high dimensional space ($\mathbb{R}^{K^j\times K^j}$). If some kernels have their distances between cluster centers smaller than those of others, their representativeness are considered less significant and the information carried by them can be transferred to other kernels. Thus, these insignificant kernels can be safely pruned without hurting the network performance. The kernel to be pruned $\mathcal{K}'^j_{m,n}$ in the jth layer is then determined as follows:

$$\arg\min_{m\in[1,C^j_{in}],\, n\in[1,C^j_{out}]} d(m,n) \tag{3}$$

where

$$d(m,n):=\left|\mu^j-\mathcal{K}^j_{m,n}\right| \tag{4}$$

Algorithm 1: Kernel Cluster Pruning

Input: Network $\mathbf{W}$, target sparsity $\mathcal{S}$, pruning epochs $\mathcal{E}$, training data $\mathbf{X}$.

Initialization: *Pretrained parameters* $\mathbf{W} = \{\mathbf{W}^{(j)},\ 1 \leq j \leq L\}$, *portion of kernels to be pruned* $p = \mathcal{S}/\mathcal{E}$;

for *epoch* $\leftarrow 1$ **to** $\mathcal{E}$ **do**
  Fine-tune model based on $\mathbf{X}$;
  for $j \leftarrow 1$ **to** L **do**
    Calculate cluster center μ^j based on Eq. 2;
    Remove kernels in the set D' of Eq. 6;
    Update network parameter $\mathbf{W}^{(j)}$ to $\mathbf{W}'^{(j)}$;
  end
  Update $p \leftarrow p + \mathcal{S}/\mathcal{E}$;
end
Obtain spare network $\mathbf{W}^*$ from $\mathbf{W}'$;

Output: Compressed model with parameter $\mathbf{W}^*$

is the distance of each kernel in the jth layer to the cluster center. To prune more than one kernel at a time, we construct a set (D) from $d\,(m,n)$ as follows:

$$D = \{x \mid x = d\,(m,n)\,,\ m \in [1, C^j_{in}],\ n \in [1, C^j_{out}]\} \quad (5)$$

and calculate different quantiles ($q(p)$) of the set in Eq. 5 to form a subset ($D' \subset D$):

$$D' = \{x' \mid x' \in D - \{d\,(m,n) \leq q(p)\}\} \quad (6)$$

p denotes the portion of kernels to be removed in that layer and $q(\cdot)$ the quantile function. Kernels in D' are those to be pruned. By gradually increasing the portion p and iteratively pruning the kernels that are closer to the cluster center, we can achieve structured sparsity. The pruned kernels are those with insignificant representativeness and the network can easily recover its performance after fine-tuning. The complete process is shown in Algo. 1.

3.3. Cost Volume Feature Distillation

Knowledge Distillation (KD) [18] refers to the technique that helps the training process of a small *student* network by having it reproduce the final output of a large *teacher* network. In our work, the *student* is the pruned network, whereas the *teacher* is the unpruned pre-trained network. Unlike the conventional KD, our method guides the pruned network to mimic the prediction of some crucial high-level building blocks of the original network; thus, the resulting performance is better than directly distilling the fine-tuned output of the teacher network.

An important trait of **stereo matching** is the computation of disparity between two views. It is crucial to fully utilize the information within cost volume for accurate stereo matching. We aim to transfer the 4D cost volume of the teacher to the student. Assuming $\mathbf{X}^{\mathcal{B}}$ to be a set of data in one mini-batch, our feature distillation (FD) loss ($\mathcal{L}_{\text{FD}}$) is expressed as follows:

$$\mathcal{L}_{\text{FD}} = \sum_{x_i \in \mathbf{X}^{\mathcal{B}}} KL\left(softmax(\frac{C_t}{T}),\ softmax(\frac{C_s}{T})\right) \quad (7)$$

where KL denotes the Kullback–Leibler divergence, whereas T denotes the temperature. In **stereo matching**, C_t and C_s are the 4D cost volumes ($height \times width \times disparity\,range \times feature\,size$) generated by the teacher and student, respectively. The final loss function applying our feature distillation is as follows:

$$\mathcal{L} = \alpha\mathcal{L}_{\text{task}} + (1-\alpha)\mathcal{L}_{\text{FD}} \quad (8)$$

where α is a constant, $\mathcal{L}_{\text{task}}$ denotes the original loss of each task, and $\mathcal{L}_{\text{FD}}$ denotes the loss described in Eq. 7.

We would like to highlight that our focus is the novel kernel pruning technique proposed in Sec. 3.2. We merely use FD as a technique to boost the accuracy after pruning. The distillation method is exclusively applied in depth estimation. There are strong correspondences of each pixel in depth estimation (the value of label is actual estimated depth values which correlate with adjacent labels), whereas in segmentation, such a connection is less significant (the value of label is merely encoding of different object classes). This indicates that our high level feature distillation method would perform better in transferring cost volume in depth estimation because cost volume carries the important correspondence information. Suppose that one is interested in segmentation distillation results, we also provide those in the supplementary materials.

4. Experiments

We first benchmark our KCP on ResNet [14] trained on CIFAR-10 [19] and ILSVRC-2012 [34] to verify the effectiveness of our method. Subsequently, for dense labeling applications, the experimental evaluations are performed on end-to-end networks. For stereo matching, we select PSMNet [2] trained on KITTI [30, 31]. HRNet [37] has been the backbone of several SOTA semantic segmentation networks. We test our pruning method on HRNet trained on Cityscapes [5] to verify the effect of pruning on semantic segmentation. Feature distillation is applied on PSMNet.

4.1. Experimental Settings

Architectures and Metrics. For ResNet, we use the default parameter settings of [14]. Data augmentation strategies are the same as in PyTorch [33]'s official examples. The pre-trained models are from official torchvision models. We report Top-1 accuracy.

For PSMNet, the pre-trained model and implementation are from original authors. The error metric is the 3-px error, which defines error pixels as those having errors of more than three pixels or 5% of disparity error. The 3-px accuracy is simply the 3-px error subtracted from 100%. We calculate the 3-px error directly using the official code implementation of the author. We prune for 300 epochs and follow the training setting of [2].

For HRNet, we use the "HRNetV2-W48" configuration with no multi-scale and flipping. The implementation and training settings follows [36]. Following the convention of semantic segmentation, the mean of class-wise intersection over union (mIoU) is adopted as the evaluation metric.

Datasets. The CIFAR-10 dataset contains 50,000 training images and 10,000 testing images. The images are 32×32 color images in 10 classes. ILSVRC-2012 contains 1.28 million training images and 50 k validation images of 1,000 classes.

The KITTI2015 stereo is a real world dataset of street views. It consists of 200 training stereo color image pairs with sparse ground-truth disparities obtained from LiDAR and 200 testing images. The training images are further divided into training set (160) and validation set (40). Image size is 376×1240.

The Cityscapes dataset contains 5,000 finely annotated color images of street views. The images are divided as follows: 2,975 for training, 500 for validation, and 1,525 for testing. There are 30 classes, but only 19 classes are used for evaluation. Image size is 1024×2048.

4.2. KCP on Classification

Table 1. Pruning results of ResNet-56/110 on CIFAR-10. The Acc. ↓ is the accuracy drop between pruned model and baseline. Negative Acc. ↓ indicates that the accuracy is better than the baseline.

Depth	Method	Baseline	Acc.	Acc.↓	Sparsity	FLOPs
56	PFEC [23]	93.04	91.31	1.73	13.70%	9.09E7
	FPGM [17]	93.59	92.93	0.66	38.71%	5.94E7
	LPFC [15]	93.59	93.34	0.25	-	5.91E7
	Ours	93.69	93.37	0.32	39.58%	**5.27E7**
	FPGM [17]	93.59	92.04	1.55	59.13%	3.53E7
	Ours	93.69	93.31	**0.38**	**60.54%**	**2.76E7**
110	PFEC [23]	93.53	92.94	0.59	32.40%	1.55E8
	FPGM [17]	93.68	93.73	-0.05	38.73%	1.21E8
	LPFC [15]	93.68	93.79	-0.11	-	1.01E8
	Ours	93.87	93.79	0.08	50.74%	**9.20E7**
	FPGM [17]	93.68	93.24	0.44	59.16%	7.19E7
	Ours	93.87	93.66	**0.21**	**59.64%**	**7.10E7**

We compare our method with other filter-level pruning works. For the CIFAR-10 dataset, we report accuracy on ResNet-56 and 110 trained from scratch. The pruning epoch $\mathcal{E}$ is set to 200 and no additional fine-tune epochs are added. As shown in Table 1, our method achieves SOTA results. Our method can achieve comparable accuracy with previous works at low sparsity, whereas at high sparsity, our method can provide more sparsity (FLOPs reduction) with less accuracy drop.

For the ILSVRC-2012 dataset, we test pre-trained ResNet-18, 34, 50, and 101. The pruning epoch $\mathcal{E}$ is set to 100. Table 2 shows that our method outperforms those of other works. For approximately 30% sparsity, our method can even achieve accuracy higher than the baseline. At higher sparsity, KCP also achieves better accuracy with larger FLOPs reduction. For example, for ResNet-50, our method achieves similar FLOPs reduction with FPGM [17] and DMC [9], but our model exceeds theirs by 1.45% and 0.93% on the accuracy, respectively (See row 8 of ResNet-50 in Table 2). Moreover, for pruning pretrained ResNet-101, our method can reduce 41.7% of FLOPs and even gain 0.55% accuracy. A similar observation can also be found on ResNet-18 and ResNet-34. KCP utilizes a smaller target for pruning and considers the relationship between different kernels, which is the reason for its success over other previous works. Proved by all these results, our method achieves SOTA results.

4.3. KCP on Depth Estimation

Recent works have shown that depth estimation from stereo pair images can be formulated as a supervised learning task resolved by neural networks. Successful depth estimation cannot be achieved without exploiting context information to find correspondence. The use of stacked 3D convolution can refine the 4D cost volume and thus provide better results. However, the 3D convolution inevitably increases the parameter and computation complexity, which enables our kernel pruning method to leverage the redundancy. Table 3 summarizes the pruning result of PSMNet [2] accuracy drop under different sparsity configurations. Our method can achieve 77% of FLOPs reduction with less than 1% of accuracy drop (P. 60.13%) and can even reduce more than 90% of FLOPs with only 3% of accuracy loss (P. 81.95%).

The predicted depth maps from KITTI2015 dataset of different sparsity models are displayed in Fig. 3. We also provide the difference map between ground truth and predictions. The darker region indicates more difference. Because the ground truths of the KITTI2015 dataset are sparse, we follow the convention by only considering the valid points during the difference map generation. Notably, the depth maps predicted by pruned models are almost identical to those of the unpruned original model. Only some edge areas are predicted differently from the ground truth. These results demonstrate that our method can effectively compress the stereo matching model with negligible performance loss.

4.4. KCP on Semantic Segmentation

High-Resolution network augments the high-resolution representations output from different stages by aggregating the (upsampled) representations from all the parallel con-

Table 2. Pruning results of ResNet on ILSVRC-2012. The ↓ has the same meaning as in Table 1.

Depth	Method	Baseline Top1/ 5	Top1(%)	Top1↓(%)	Top5(%)	Top5↓(%)	Sparsity	FLOPs	FLOPs↓
18	FPGM [17]	70.28 / 89.63	68.41	1.87	88.48	1.15	28.10%	1.08E9	41.8%
	DMCP [10]	70.10 / -	69.20	0.9	-	-	-	1.04E9	43.8%
	Ours	69.76 / 89.08	70.10	**-0.34**	89.45	**-0.33**	27.2%	1.06E9	42.8%
	FPGM [17]	70.28 / 89.63	62.23	8.05	84.31	5.32	56.34%	5.53E8	70.2%
	Ours	69.76 / 89.08	63.79	**5.97**	85.56	**3.52**	55.86%	**4.70E8**	**75.1%**
34	PFEC [23]	73.23 / -	72.17	1.06	-	-	10.80%	2.76E9	24.2%
	FPGM [17]	73.92 / 91.62	72.63	1.29	91.08	0.54	28.89%	2.18E9	41.1%
	DMC [9]	73.30 / 91.42	72.57	0.73	91.11	0.31	-	2.10E9	43.4%
	Ours	73.31 / 91.42	73.46	**-0.15**	91.56	**-0.14**	**29.25%**	**2.04E9**	**44.9%**
	FPGM [17]	73.92 / 91.62	67.69	6.23	87.97	3.65	57.94%	1.08E9	70.9%
	Ours	73.31 / 91.42	69.19	**4.12**	88.89	**2.53**	57.93%	**9.19E8**	**75.2%**
50	DMCP [10]	76.60 / -	77.00	-0.40	-	-	-	2.80E9	27.4%
	SFP [16]	76.15 / 92.87	62.14	14.01	84.60	8.27	30.00%	2.24E9	41.8%
	HRank [24]	76.15 / 92.87	74.98	1.17	92.33	0.51	36.67%	2.30E9	43.8%
	FPGM [17]	76.15 / 92.87	75.59	0.56	92.23	0.24	24.26%	2.23E9	42.2%
	Ours	76.13 / 92.87	76.61	**-0.48**	93.21	**-0.34**	31.19%	**2.18E9**	43.5%
	FPGM [17]	76.15 / 92.87	74.83	1.32	92.32	0.55	32.37%	1.79E9	53.5%
	DMC [9]	76.15 / 92.87	75.35	0.80	92.49	0.38	-	1.74E9	55.0%
	Ours	76.13 / 92.87	76.26	**-0.13**	93.04	**-0.17**	**37.06%**	**1.77E9**	**54.1%**
	LFPC [15]	76.15 / 92.87	74.46	1.69	92.04	0.83	-	1.51E9	60.8%
	Ours	76.13 / 92.87	74.80	**1.33**	92.19	**0.68**	46.11%	**1.30E9**	**66.3%**
101	FPGM [17]	77.37 / 93.56	77.32	0.05	93.56	0.00	26.63%	4.37E9	42.2%
	Ours	77.37 / 93.56	77.92	**-0.55**	93.92	**-0.36**	28.48%	4.41E9	41.7%

Table 3. Accuracy and FLOPs of different sparsity configuration of PSMNet on KITTI2015. P.30.52% means that the model is pruned to 30.52% sparsity. The FLOPs are calculated under the assumption that the network runs at full HD resolution and 60 fps.

	3-px Acc.	3-px Acc.↓	FLOPs	FLOPs ↓
Original	99.10%	-	2.43E14	-
P. 30.52%	98.81%	0.29%	1.32E14	45.68%
P. 51.46%	98.56%	0.54%	7.47E13	69.26%
P. 60.13%	98.50%	0.61%	5.54E13	77.20%
P. 70.56%	97.59%	1.52%	3.57E13	85.31%
P. 81.95%	96.07%	3.06%	1.85E13	92.39%

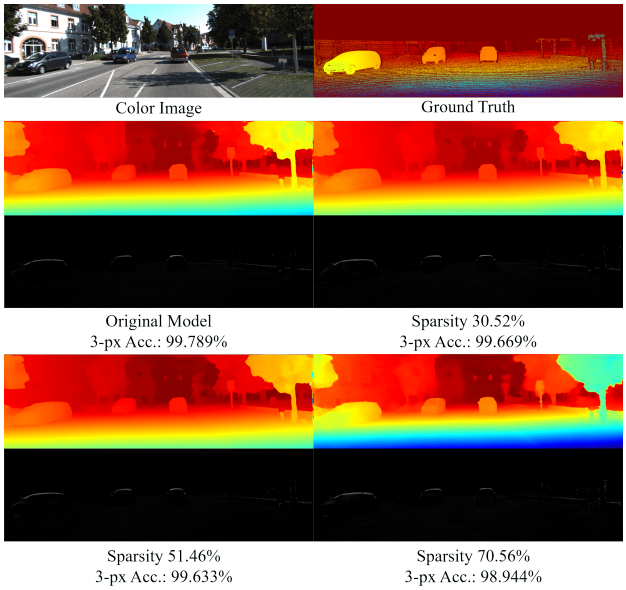

Figure 3. Results of our method of PSMNet on KITTI2015. We show the predicted depth map and the difference with ground truth (brightness indicates difference) from unpruned original model, 30% sparsity, 50% sparsity and 70% sparsity. The 3-px accuracies are listed below each predicted result. The model performance is well preserved both qualitatively and quantitatively.

volutions. The different resolution blocks comprise convolutional layers with a large receptive field, which contains significant redundancy for our kernel pruning method to optimize. Fig. 4 illustrates the change of validation mean IoU of HRNet [37] pruned with different $\mathcal{S}$. Notably, at sparsity of less than 30%, the performance of the pruned model can even exceed that of the original model. The mean of IoU degradation is within 2.7% even at 70% sparsity. Table 4 summarizes the mIoU drop and resulting FLOPs of the pruned models. Based on the result of HRNet pruning, our method is proven to be effective on semantic segmentation, again we can reduce 35% of FLOPs with 0.38% of performance gain, and can also reduce 74% of FLOPs with less than 1% of performance loss. The visualization results of the predicted class label of Cityscapes validation

Table 4. Mean IoU and FLOPs of different sparsity of HRNet on Cityscapes. Negative meanIoU↓ means that the meanIou is higher than the original unpruned model. P.20.35% denotes that the model is pruned to 20.35% sparsity. The FLOPs are calculated with the network taking an input of size 512×1024.

	mIoU	mIoU↓(%)	FLOPs	FLOPs↓(%)
Original	0.8090	-	1.74E11	-
P. 20.35%	0.8121	**-0.383**	1.13E11	35.06%
P. 29.95%	0.8096	**-0.074**	8.78E10	49.54%
P. 50.31%	0.8040	0.618	4.39E10	74.77%
P. 59.90%	0.7924	2.052	2.74E10	84.25%
P. 69.49%	0.7873	2.682	1.35E10	92.24%

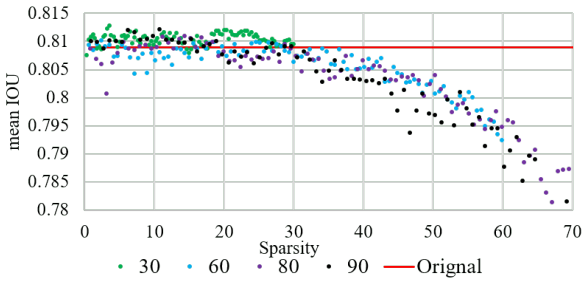

Figure 4. Validation mean IoU of HRNet pruned with different $\mathcal{S}$.

images are shown in Fig. 5. To highlight the difference, we calculate the difference map between predicted class labels and ground truth on the 19 classes used for evaluation. The black region indicates identical class labels, whereas the white region indicates different class labels with ground truth. Other classes appearing in the ground truth, but are ignored in evaluation are not shown on the difference map. The difference maps of the pruned models in Fig. 5 are almost identical to those of the original model. The results demonstrate that our method can remove redundant weights without hurting the network performance.

4.5. Ablation Study

We conduct experiments to perform ablation studies in this section. The effect of the proposed cost volume feature distillation will be presented. The adversarial criteria is also added to further validate the effectiveness of our method.

Feature Distillation. Variants of feature distillation are applied to demonstrate the appropriate distillation target for increasing model performance. Corresponding results of PSMNet are presented in Fig. 6. For the network being pruned, "No FD" does not apply feature distillation, "Cost Volume FD" utilizes 4D cost volume as the distillation target as described in Eq. 7, and "Disparity FD" uses final disparity output as the target for the conventional knowledge distillation. We set α to 0.9 and T to 15. Results shown in Fig. 6 well demonstrate that by encouraging the pruned network to mimic the 4D cost volume construction of the original network (Cost Volume FD), the accuracy after pruning can be better maintained. Distilling the final output even results in worse accuracy than not applying distillation. We believe that the correct choice of crucial feature is the key to better model performance. Because the construction of cost volume plays a critical role in stereo matching, distilling the feature of cost volume provides the best results.

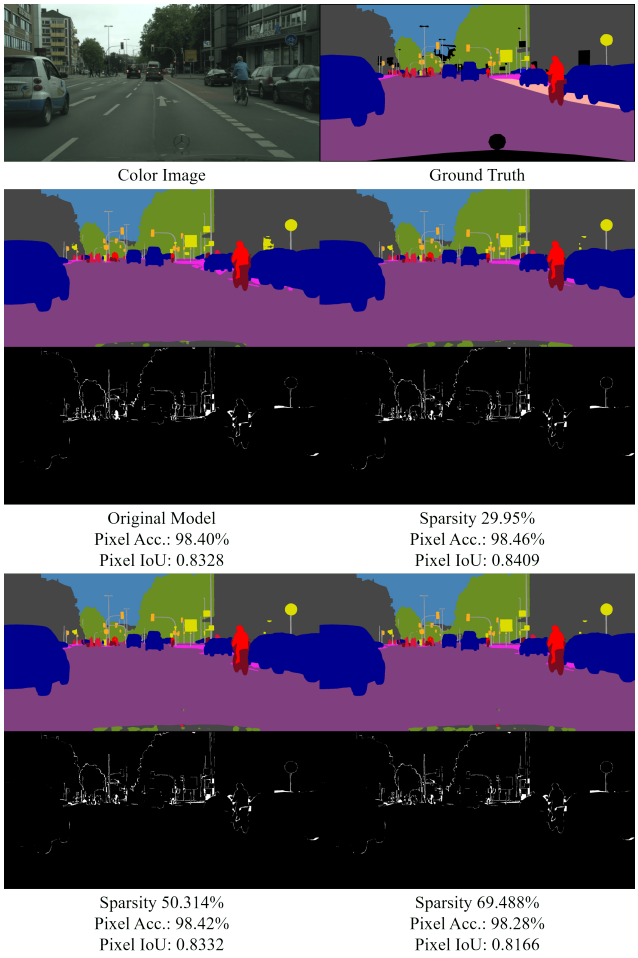

Figure 5. Results of our method of HRNet on Cityscapes. The predicted class label and the difference between ground truth are shown, (brightness indicates difference) including original model, 30% sparsity, 50% sparsity and 70% sparsity. The pixel accuracy and pixel IoU are listed below each figure.

Adversarial Criteria. We apply adversarial criteria to further validate the effectiveness of our method. Adversarial criteria means that in contrast to Eq. 3 in which we prune the kernels that are the **closest** to the cluster center, we instead remove the kernels that are the **farthest** from the cluster center. Table 5 shows that even only replacing 10% of our method into adversarial criteria significantly degrades the accuracy (10% Adv.). The accuracy becomes worse when the percentage of adversarial criteria increases,

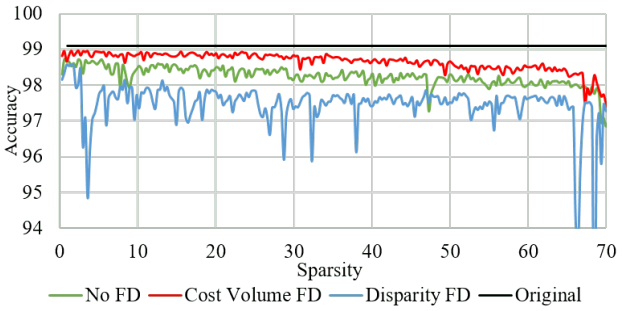

Figure 6. Variants of feature distillation for PSMNet pruning.

Table 5. Analysis of adversarial criteria. Models are ResNet-56 and ResNet-110. $\mathcal{S}$ is set to 30. 10% Adv. denotes that 10% of the layers are randomly chosen to be pruned using adversarial criteria.

	ResNet-56		ResNet-110	
Config.	Acc.	Sparsity	Acc.	Sparsity
Non Adv.	93.62%	29.69%	94.05%	29.80%
10% Adv.	74.85%	30.14%	76.88%	29.47%
20% Adv.	51.96%	28.13%	36.75%	26.99%
40% Adv.	14.35%	29.72%	13.28%	27.23%

indicating that our method can truly select the kernels that are the least representational in each layer; evidence of the effectiveness of our proposed method.

4.6. Further Exploration

Table 6. Mean PSNR and SSIM of unpruned and pruned SRGAN generator on recovering DIV2K validation images.

	PSNR (dB)	SSIM
Finetune	28.1009	0.7968
Prune 37.16%	27.6507	0.7937
Prune 53.16%	27.7121	0.7887

To further validate our KCP, we apply pruning on SRGAN [21], which recovers a single image resolution using a generative adversarial network (GAN). We are curious about how pruning affects the image generation of GAN and the perceptual difference induced by pruning on restored high-resolution images.

The SRGAN is first trained on color images sampled from VOC2012 [7] dataset for 100 epochs and then finetuned on DIV2K [1] dataset. We only prune the generator and keep the discriminator updated, because during inference, only the generator is used to recover low resolution images. The experiments are performed with a scale factor of $4\times$ between low- and high-resolution images. We report two classical super-resolution performance measurements

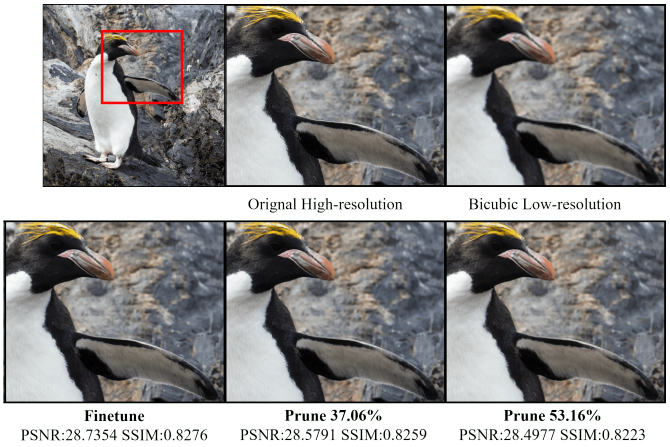

Figure 7. Corresponding high resolution image, bicubic-upsampled low resolution image, and reconstruction results from different sparsity SRGAN generators. The image is from DIV2K validation set. Results are zoomed-in from the red block.

PSNR (dB) and SSIM for comparison.

Table 6 shows the results. The pruned generator can achieve similar PSNR and SSIM to the unpruned model. In terms of human perception, Fig. 7 demonstrates that the pruned generator still reconstructs the details quite well. Little difference can be observed by human perception. This means that our KCP successfully selected the low representative kernels and the success on GAN pruning paves the way for explainable pruning. In the future, the interaction between specific parameters and human perception may be observed through more experiments on GAN pruning.

5. Conclusions

In this paper, we present a novel structured pruning method KCP, which explores the representativeness of parameters by clustering filter kernels, to elucidate dense labeling neural network pruning. We demonstrate that by removing the kernels that are closest to the cluster center in each layer, we can achieve high FLOPs reduction with negligible performance loss on dense labeling application models. In addition, we propose cost volume feature distillation to boost the accuracy of depth estimation network after pruning. Extensive experiments on various CNNs demonstrate the effectiveness of KCP in reducing model size and computation overhead while showing little compromise on the model performance. This paves the way for efficient model deployment of those real-time dense labeling algorithms. Notably, KCP achieves new SOTA results for ResNet on both CIFAR-10 and ImageNet datasets. We also explore the perception difference induced by pruning through applying KCP on GAN. We believe that in the future, generative model pruning may help researchers to develop human perception based pruning.

References

[1] E. Agustsson and R. Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh. (CVPRw)*, pages 1122–1131, 2017. 8

[2] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 5410–5418, 2018. 1, 4, 5

[3] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)*, 40(4):834–848, 2017. 1

[4] Ting-Wu Chin, Ruizhou Ding, Cha Zhang, and Diana Marculescu. Towards efficient model compression via learned global ranking. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 1518–1528, 2020. 3

[5] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2016. 1, 4

[6] Yann Le Cun, John S. Denker, and Sara A. Solla. *Optimal Brain Damage*, page 598–605. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1990. 2

[7] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *Int. J. Comput. Vis. (IJCV)*, 111(1):98–136, 2015. 8

[8] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*, 2018. 2

[9] Shangqian Gao, Feihu Huang, Jian Pei, and Heng Huang. Discrete model compression with resource constraint for deep neural networks. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 1899–1908, 2020. 3, 5, 6

[10] Shaopeng Guo, Yujie Wang, Quanquan Li, and Junjie Yan. Dmcp: Differentiable markov channel pruning for neural networks. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 1539–1547, 2020. 2, 3, 6

[11] Song Han, Huizi Mao, and William J. Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding, 2015. 2

[12] Song Han, Jeff Pool, John Tran, and William J. Dally. Learning both weights and connections for efficient neural networks. In *Adv. Neural Inform. Process. Syst. (NIPS)*, page 1135–1143, Cambridge, MA, USA, 2015. MIT Press. 2

[13] Babak Hassibi and David G Stork. Second order derivatives for network pruning: Optimal brain surgeon. In *Advances in neural information processing systems*, pages 164–171, 1993. 2

[14] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 770–778, June 2016. 1, 3, 4

[15] Yang He, Yuhang Ding, Ping Liu, Linchao Zhu, Hanwang Zhang, and Yi Yang. Learning filter pruning criteria for deep convolutional neural networks acceleration. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 2009–2018, 2020. 1, 2, 5, 6, 11

[16] Yang He, Guoliang Kang, Xuanyi Dong, Yanwei Fu, and Yi Yang. Soft filter pruning for accelerating deep convolutional neural networks. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2234–2240, 2018. 2, 6

[17] Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2019. 1, 2, 5, 6

[18] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 4

[19] Alex Krizhevsky et al. Learning multiple layers of features from tiny images. 2009. 2, 3, 4, 11

[20] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 1

[21] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 105–114, 2017. 8

[22] Carl Lemaire, Andrew Achkar, and Pierre-Marc Jodoin. Structured pruning of neural networks with budget-aware regularization. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 9108–9116, 2019. 2

[23] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710*, 2016. 2, 5, 6

[24] Mingbao Lin, Rongrong Ji, Yan Wang, Yichen Zhang, Baochang Zhang, Yonghong Tian, and Ling Shao. Hrank: Filter pruning using high-rank feature map. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 1529–1538, 2020. 1, 2, 6

[25] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *Int. Conf. Comput. Vis. (ICCV)*, pages 2736–2744, 2017. 2

[26] Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning. *arXiv preprint arXiv:1810.05270*, 2018. 2

[27] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 3431–3440, 2015. 1

[28] Jian-Hao Luo and Jianxin Wu. Neural network pruning with residual-connections and limited-data. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 1458–1467, 2020. 3

[29] Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. Thinet: A filter level pruning method for deep neural network compression. In *Int. Conf. Comput. Vis. (ICCV)*, pages 5058–5066, 2017. 2

[30] Moritz Menze, Christian Heipke, and Andreas Geiger. Joint 3d estimation of vehicles and scene flow. In *ISPRS Workshop on Image Sequence Analysis (ISA)*, 2015. 1, 4

[31] Moritz Menze, Christian Heipke, and Andreas Geiger. Object scene flow. *ISPRS Journal of Photogrammetry and Remote Sensing (JPRS)*, 2018. 1, 4

[32] Pavlo Molchanov, Arun Mallya, Stephen Tyree, Iuri Frosio, and Jan Kautz. Importance estimation for neural network pruning. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 11264–11272, 2019. 2

[33] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 4

[34] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis. (IJCV)*, 115, 09 2014. 2, 3, 4

[35] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1

[36] Ke Sun, Yang Zhao, Borui Jiang, Tianheng Cheng, Bin Xiao, Dong Liu, Yadong Mu, Xinggang Wang, Wenyu Liu, and Jingdong Wang. High-resolution representations for labeling pixels and regions, 2019. 5

[37] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep high-resolution representation learning for visual recognition. *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)*, 2019. 1, 4, 6

[38] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. In *Advances in neural information processing systems*, pages 2074–2082, 2016. 2

[39] Jianbo Ye, Xin Lu, Zhe L. Lin, and James Zijun Wang. Rethinking the smaller-norm-less-informative assumption in channel pruning of convolution layers. *ArXiv*, abs/1802.00124, 2018. 2

[40] Po-Hsiang Yu, Sih-Sian Wu, Jan P. Klopp, Liang-Gee Chen, and Shao-Yi Chien. Joint pruning & quantization for extremely sparse neural networks. *arXiv preprint arXiv:2010.01892*, 2020. 1, 3

[41] Jure Žbontar and Yann LeCun. Stereo matching by training a convolutional neural network to compare image patches. *The journal of machine learning research*, 17(1):2287–2318, 2016. 1

A. ResNet on CIFAR-100

CIFAR-100 is a dataset just like the CIFAR-10 [19], except that it has 100 classes containing 600 images each. There are 500 training images and 100 testing images per class. The experimental setting of ResNet on CIFAR-100 follows the setting mentioned in Sec. 4.1. Since there are few methods having experiments on CIFAR-100, we only list LFPC [15] for comparison. The result of pruning ResNet-56 on CIFAR-100 is presented in Table 7. When achieving slightly better FLOPs reduction, our KCP can result in better accuracy after pruning. In fact, the accuracy of KCP-pruned ResNet-56 even increases by 0.23%. This further proves the effectiveness of our method.

Additional pruning results for different depth ResNet on CIFAR-100 are listed in Table 8.

Table 7. Comparison of pruning ResNet-56 on CIFAR-100.

Depth	Method	Acc.↓(%)	FLOPs	FLOPs ↓
56	LFPC [15]	0.58	6.08E+07	51.60%
	Ours	-0.23	5.95E+07	52.57%

B. More Ablation Study

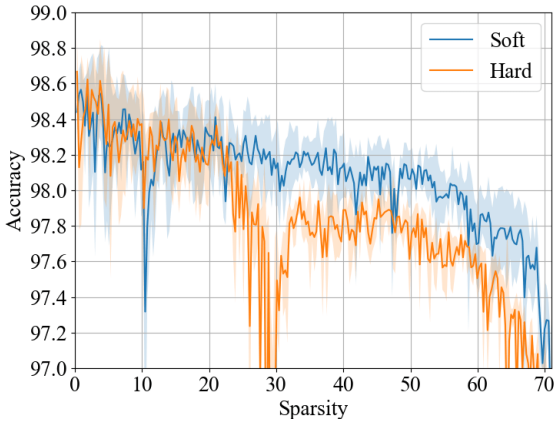

Figure 8. The influence of different fine-tuning strategy. Line represents the average over 3 experiment and the shaded area are the upper and lower bound of the standard deviation.

Fine-tune Strategy and Kernel Normalization. For brevity, we report the ablative results of the PSMNet on these two variant categories. Similar observation can be found in other network and dataset that we used in our work. Two common fine-tune strategies, soft and hard, are examined with our kernel cluster pruning. Soft means the removed kernel can be recovered with updated value throughout the pruning process, while hard means the removed weight remains zeros whenever it is pruned. As previous

Table 8. Results of KCP processed ResNet on CIFAR-100.

Baseline (Original Unpruned Models)		
Depth	Acc.	FLOPs
20	68.23	4.06E+07
32	69.62	6.89E+07
56	71.27	1.25E+08
86	72.76	1.96E+08
110	73.31	2.53E+08

KCP Pruned Models				
Depth	Acc. ↓(%)	Sparsity	FLOPs	FLOPs ↓
20	1.07	34.20%	1.98E+07	51.22%
	2.49	38.66%	1.76E+07	56.60%
	3.45	51.58%	1.20E+07	70.47%
32	-0.06	34.76%	3.32E+07	51.86%
	0.89	43.42%	2.62E+07	86.64%
	1.96	52.10%	2.00E+07	70.91%
56	-0.32	26.23%	7.43E+07	40.80%
	-0.23	35.38%	5.95E+07	52.57%
	0.42	43.96%	4.71E+07	62.49%
	0.89	53.33%	3.51E+07	72.05%
	2.89	60.84%	2.67E+07	78.76%
86	-0.29	27.49%	1.13E+08	42.36%
	0.65	36.44%	9.09E+07	53.71%
	0.76	44.08%	7.37E+07	62.46%
	1.52	52.89%	5.60E+07	71.46%
	1.21	61.32%	4.12E+07	79.00%
110	0.43	28.43%	1.73E+08	31.75%
	0.64	35.52%	1.20E+08	52.63%
	1.06	45.37%	9.14E+07	63.86%
	1.44	52.36%	7.34E+07	70.97%
	1.92	61.08%	5.36E+07	78.82%

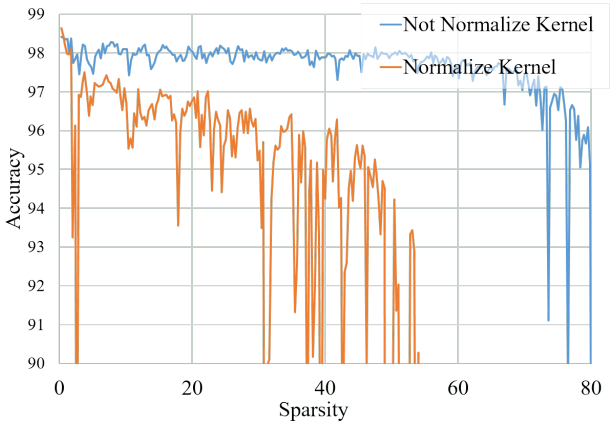

Figure 9. The influence of normalizing the kernel in each layer before performing clustering.

mentioned in Sec 3.2, we adopted soft fine-tuning in our method. Fig. 8 validates that soft fine-tuning stabilizes the pruning process and results in better accuracy. Soft fine-tuning recovers pruned kernels, which keep the size of network representation space in terms of kernel numbers. The full space enables our method to accurately identify the lowest symbolic kernels in each iteration. We also test the effect of normalizing the kernels in each layer before clustering. Result in Fig. 9 suggests that normalization hurts the performance. Since our method relies on the interaction between kernels in each layer, normalizing kernels breaks the primitive relation, thus decreases the accuracy.

C. Feature Distillation of HRNet

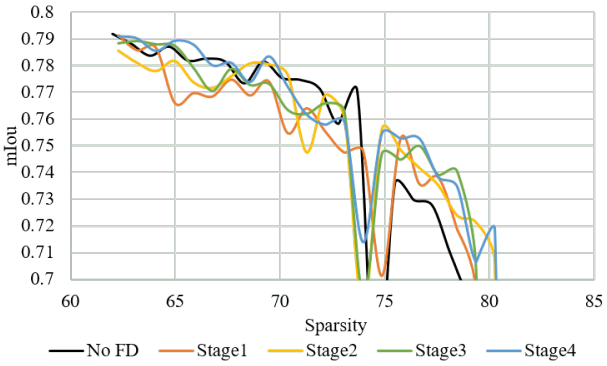

Figure 10. Variants of feature distillation for HRNet

As mentioned in Sec. 3.3, features distillation of depth estimation cannot be equally applied on segmentation. Cost volume is a significant target for stereo depth estimation to transfer, while there is no such equivalent in segmentation. In low sparsity range, distilling the intermediate stage output results in worse accuracy. Fig. 10 shows the distillation result of sparsity above 60%. Larger stage denotes that the feature used for distillation is extracted from the layer closer to the final output. Stage1 feature preserves more characteristic of the input, while Stage4 feature is more like the final predicted class label. The results show that despite the accuracy of model using distillation can exceed model without distillation at some point, none of the distillation curves (Stage1 to Stage4 curves) constantly achieves better accuracy than none distillation model (No FD curve) below 75% sparsity. At sparsity above 75%, the accuracy drop becomes larger (yet still within 10%) and unstable. In summary, feature distillation can still help to retain the accuracy of HRNet at high sparsity range, but the effect is much less significant than cost volume feature distillation (c.f Sec 3.3) of stereo depth estimation networks.

D. Adversarial Criteria on Stereo Matching and Semantic Segmentation

We also apply adversarial criteria on PSMNet and HRNet to validate our method. The results in Table 9 and 10 are coherent with Table 5. Adversarial criteria greatly affect the performance of KCP, which well demonstrates that the proposed original KCP is an effective method in terms of finding low representational kernels then remove them.

Table 9. Results of adversarial criteria applied on PSMNet. $\mathcal{S}$ is set to 90. The results are extracted from different epoch during the training process. Non Adv. denotes original KCP. 10% Adv. denotes that 10% of the layers are randomly chosen to be pruned using adversarial criteria and vice versa.

Spar.	Non Adv.	10% Adv.	20% Adv.	40% Adv.
≈10%	98.60%	96.57%	95.93%	91.22%
≈20%	98.51%	94.46%	90.39%	75.26%
≈30%	98.37%	92.92%	72.11%	45.01%
≈40%	98.30%	79.10%	35.47%	13.87%
≈50%	98.18%	70.81%	32.66%	14.09%
≈60%	97.77%	37.17%	14.10%	-
≈70%	97.41%	59.42%	-	-

Table 10. Results of adversarial criteria applied on HRNet. $\mathcal{S}$ is set to 30. The results are extracted from different epoch during the training process. Non Adv. denotes original KCP. 2% Adv. denotes that 2% of the layers are randomly chosen to be pruned using adversarial criteria and vice versa.

Spar.	Non Adv.	2% Adv.	5% Adv.	10% Adv.
≈ 5%	0.8111	0.7828	0.6589	0.6571
≈ 10%	0.8116	0.7517	0.6317	0.4938
≈ 15%	0.8109	0.7450	0.5866	0.2456
≈ 20%	0.8121	0.7125	0.2470	0.2028
≈ 30%	0.8096	0.6058	-	-

E. More Qualitative Results

We provide more visualization results of KCP on different applications. Fig. 11 and 12 are predicted depth from KITTI2015 of PSMNet. The layout is identical with Fig 3 in the main manuscript. Fig. 13 and 14 are predicted class labels from Cityscapes of HRNet. The layout is identical with Fig 5 in the main manuscript. Fig. 15 and 16 are SRGAN reconstruction results of image from DIV2K dataset.

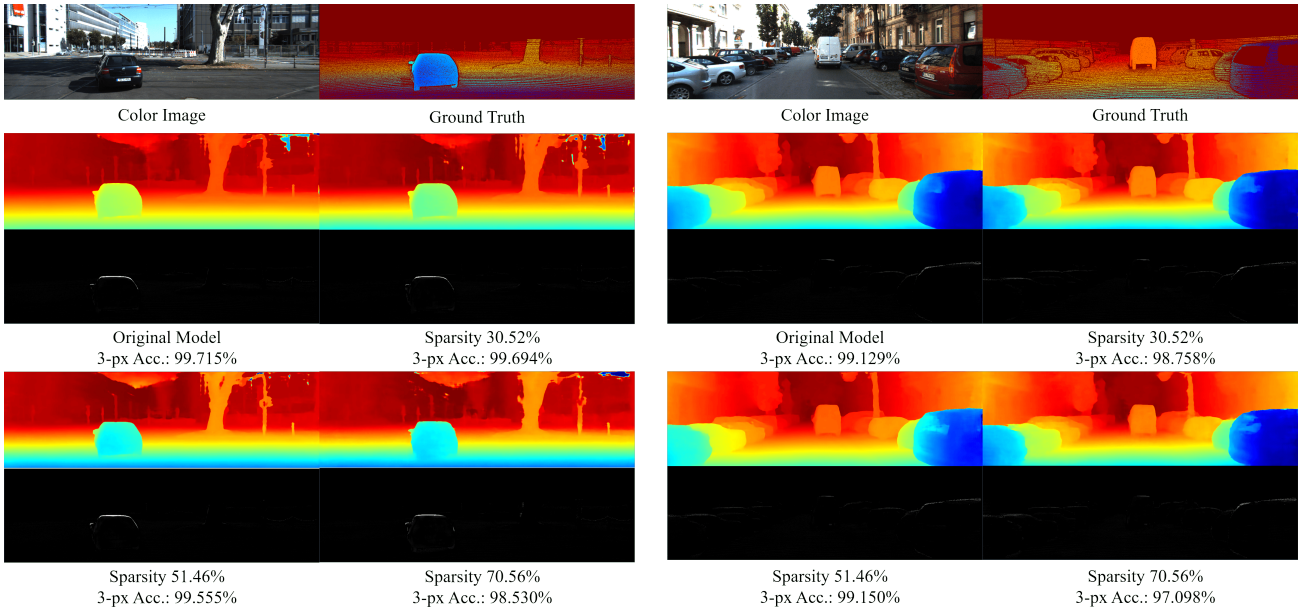

Figure 11. Results of our method of PSMNet on KITTI2015. We show the predicted depth map and the difference with ground truth (brightness indicates difference) from unpruned original model, 30% sparsity, 50% sparsity and 70% sparsity. The 3-px accuracies are listed below each predicted results.

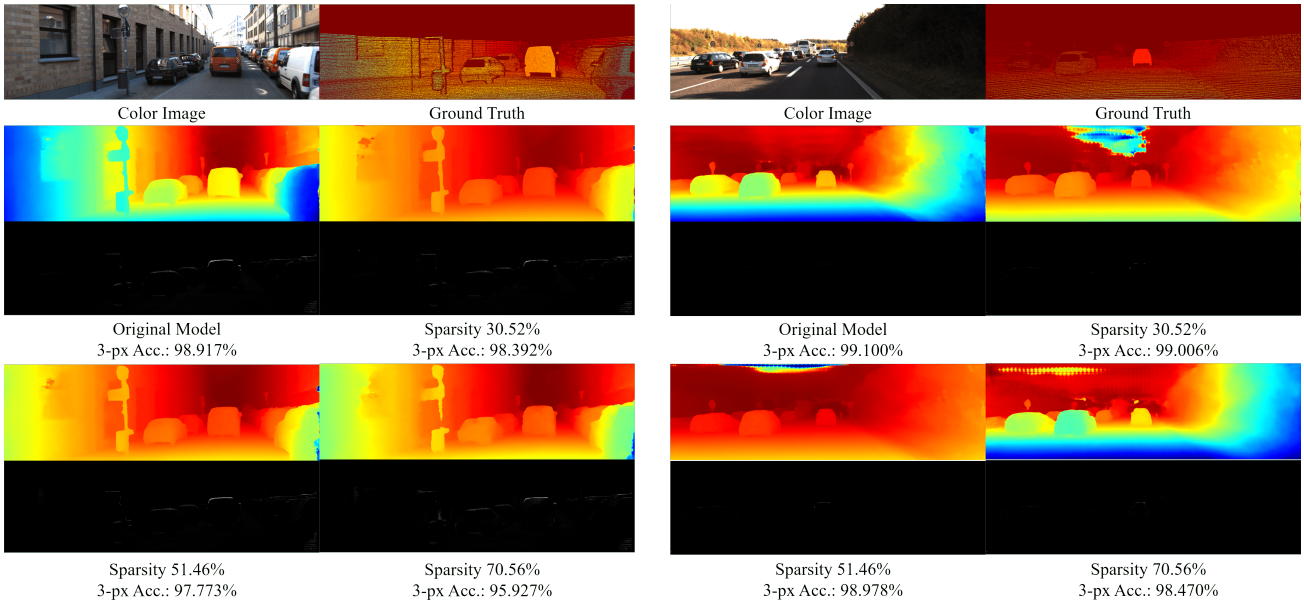

Figure 12. Results of our method of PSMNet on KITTI2015. We show the predicted depth map and the difference with ground truth (brightness indicates difference) from unpruned original model, 30% sparsity, 50% sparsity and 70% sparsity. The 3-px accuracies are listed below each predicted results.

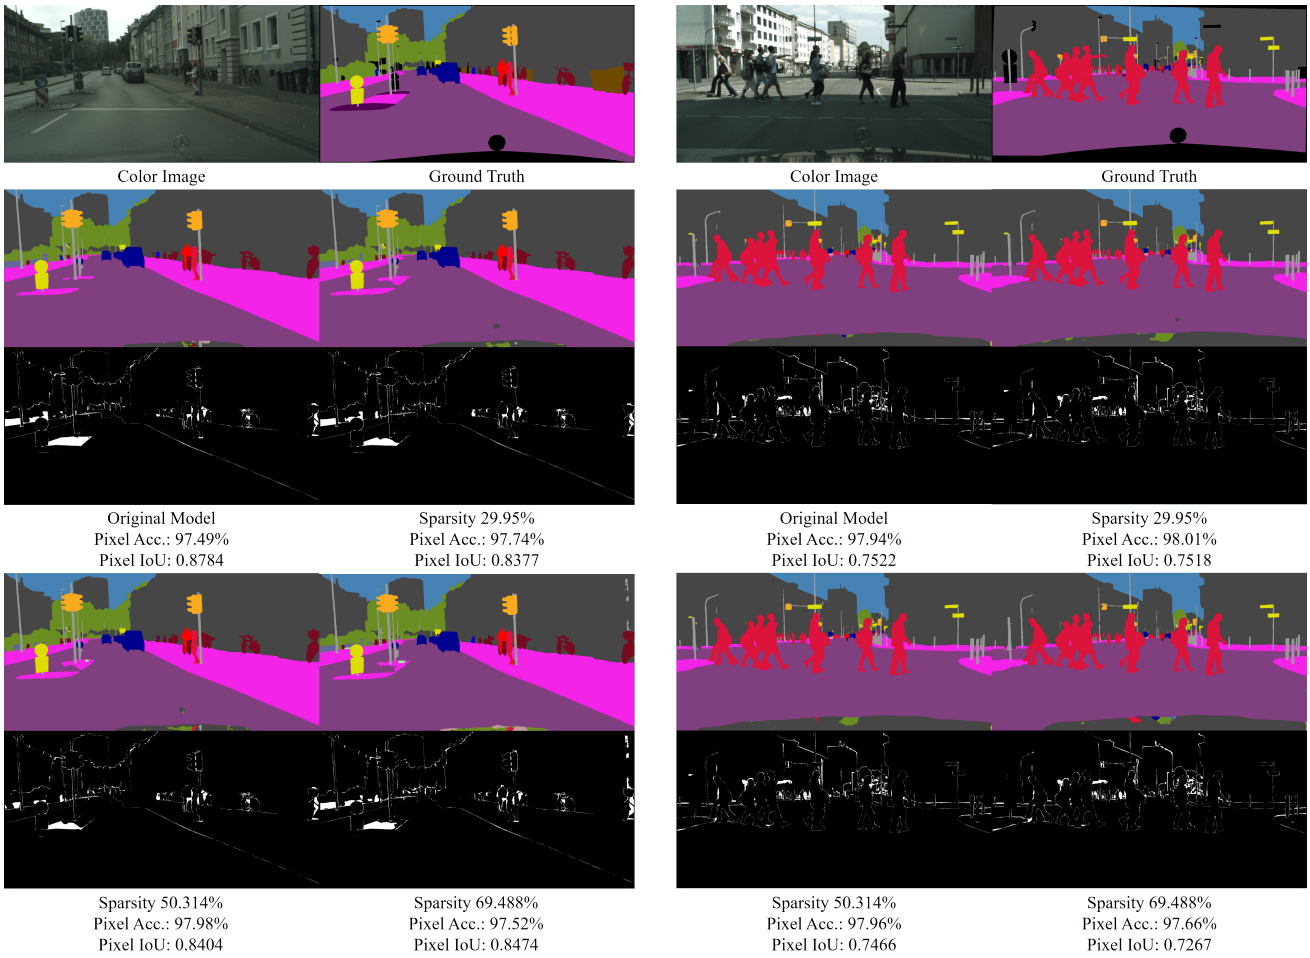

Figure 13. Results of our method of HRNet on Cityscapes. The predicted class label and the difference between ground truth are shown, (brightness indicates difference) including original model, 30% sparsity, 50% sparsity and 70% sparsity. The pixel accuracy and pixel IoU are listed below each figure.

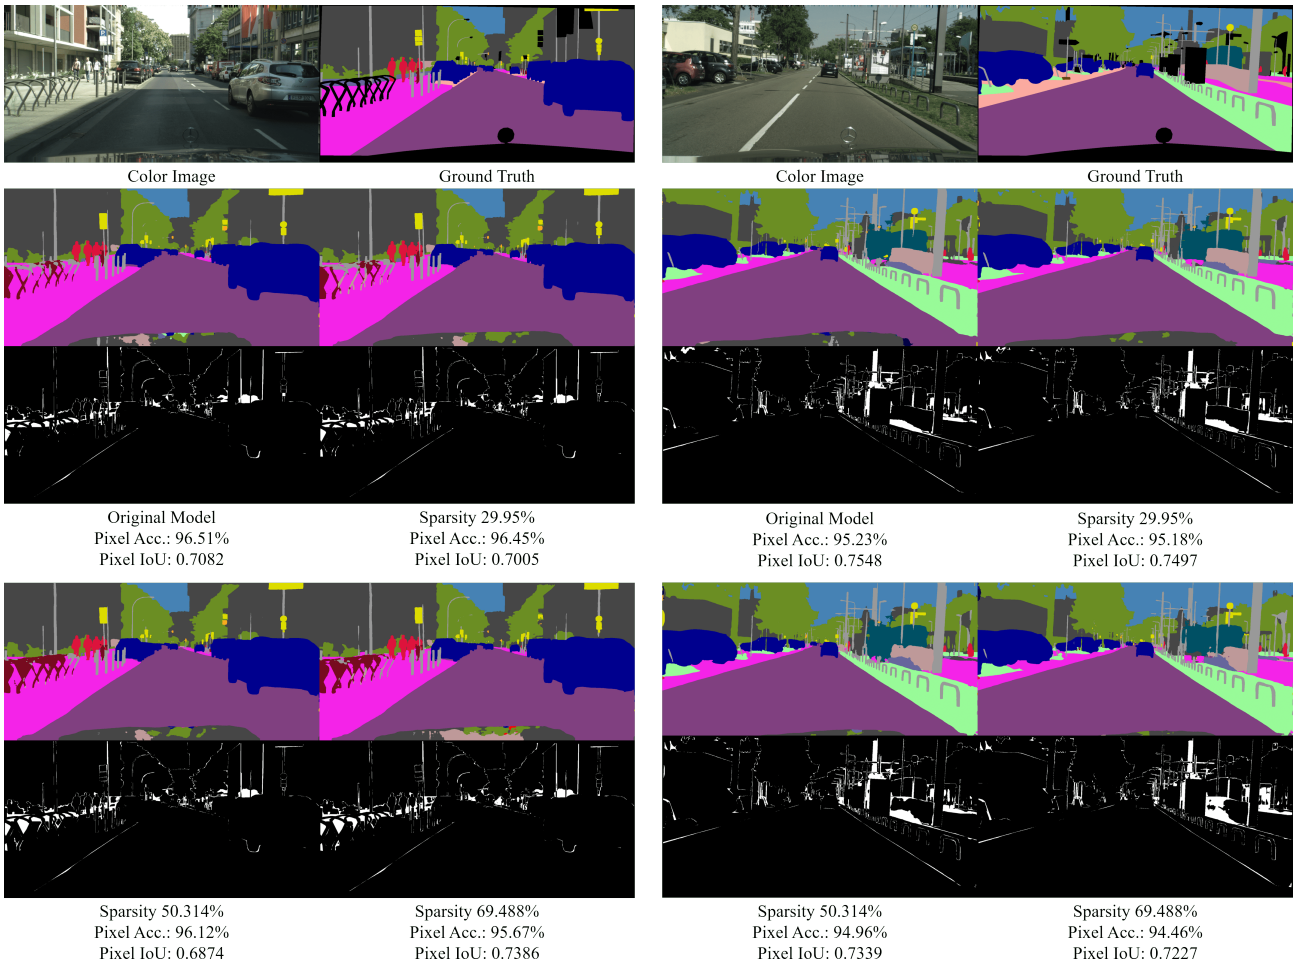

Figure 14. Results of our method of HRNet on Cityscapes. The predicted class label and the difference between ground truth are shown, (brightness indicates difference) including original model, 30% sparsity, 50% sparsity and 70% sparsity. The pixel accuracy and pixel IoU are listed below each figure.

Figure 15. Corresponding high resolution image, bicubic-upsampled low resolution image and reconstruction results from different sparsity SRGAN generators. The image is from DIV2K validation set. Results are zoomed-in from the red block.

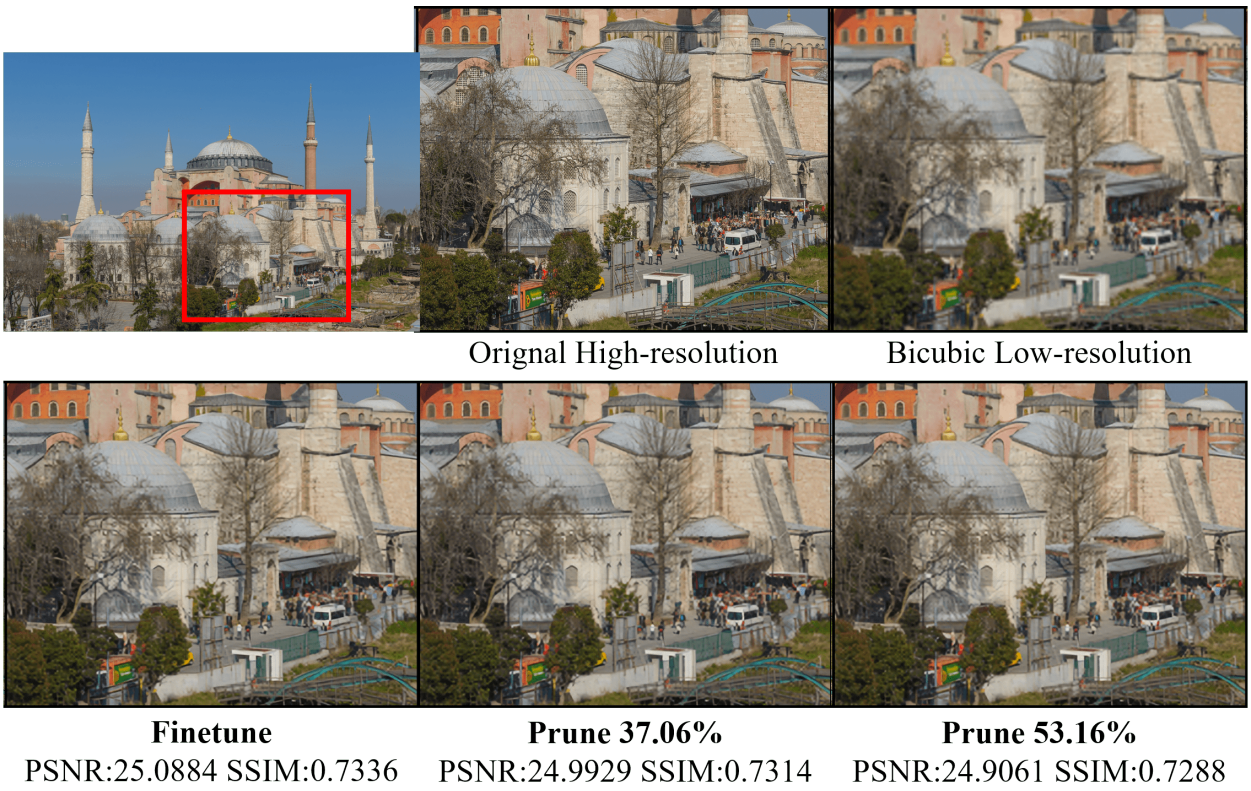

Figure 16. Corresponding high resolution image, bicubic-upsampled low resolution image and reconstruction results from different sparsity SRGAN generators. The image is from DIV2K validation set. Results are zoomed-in from the red block.