

Malleable 2.5D Convolution: Learning Receptive Fields along the Depth-axis for RGB-D Scene Parsing

Yajie Xing¹[0000-0002-1226-1529], Jingbo Wang²[0000-0001-9700-6262], and Gang Zeng¹[0000-0002-9575-4651]

¹ Key Laboratory of Machine Perception, Peking University, China
{yajie_xing,zeng}@pku.edu.cn

² The Chinese University of Hong Kong
jbwang@ie.cuhk.edu.hk

Abstract. Depth data provide geometric information that can bring progress in RGB-D scene parsing tasks. Several recent works propose RGB-D convolution operators that construct receptive fields along the depth-axis to handle 3D neighborhood relations between pixels. However, these methods pre-define depth receptive fields by hyperparameters, making them rely on parameter selection. In this paper, we propose a novel operator called malleable 2.5D convolution to learn the receptive field along the depth-axis. A malleable 2.5D convolution has one or more 2D convolution kernels. Our method assigns each pixel to one of the kernels or none of them according to their relative depth differences, and the assigning process is formulated as a differentiable form so that it can be learnt by gradient descent. The proposed operator runs on standard 2D feature maps and can be seamlessly incorporated into pre-trained CNNs. We conduct extensive experiments on two challenging RGB-D semantic segmentation dataset NYUDv2 and Cityscapes to validate the effectiveness and the generalization ability of our method.

Keywords: RGB-D Scene Parsing, Geometry in CNN, Malleable 2.5D Convolution

1 Introduction

Recent progresses[25, 2, 3] in CNN have achieved great success in scene parsing tasks such as semantic segmentation. Depth data provide geometric information that is not captured by the color channels and therefore can assist feature extraction and improve segmentation performance. With the availability of commercial RGB-D sensors such as Kinect, there comes an increasing interest in exploiting the additional depth data and incorporating geometric information into CNNs.

Many works[25, 17, 10, 6, 16] take RGB images and depth maps or HHA[9] encodings as two separate inputs and adopt two-stream style networks to process them. These methods only take depth information as features and keep the fixed geometric structures of 2D CNN, which neglects the available 3D geometric

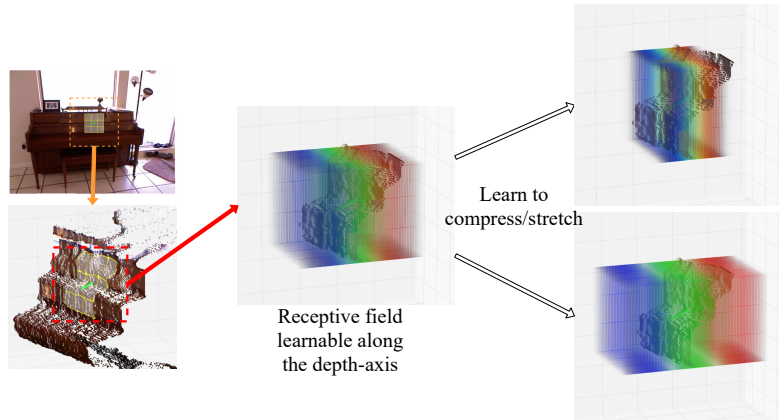


Fig. 1. Illustration of the malleable 2.5D convolution with 3 kernels. Best view in color. The blue, green, and red colors respectively represent the receptive field of each kernel. The malleable 2.5D convolution arranges its kernels sequentially along the depth-axis and adopts differentiable functions to construct soft depth receptive fields for its kernels. It samples pixels on the 2D plane, the same as standard 2D convolution. And it assigns pixels to its kernels according to depth data and its depth receptive fields. During the training process, the depth receptive fields can be learnt to compress or stretch automatically according to the dataset, making the convolution "malleable"

relations between pixels. The most direct method to leverage the 3D relations between pixels is to project RGB-D data into 3D pointclouds[23] or volumes[29, 39], and exploit 3D networks to handle geometry. However, 3D networks are computationally more expensive than 2D CNN and cannot benefit from prevalent ImageNet-pretrained models.

Recently, some works turn to introduce geometric information into the 2D convolution operators. Depth-aware CNN[31] augments the standard convolution with a depth similarity term. It applies masks to suppress the contribution of pixels whose depths are different from the center of the kernel. Consequently, it constructs a soft receptive field along the depth-axis where more distant pixels are partially occluded by the mask. 2.5D convolution[34] moves a step forward to utilizing more kernels to capture richer geometric relations. It builds a grid receptive field in 3D space, assigns each pixel to one of the kernels according to the relative depth differences with the center of the kernel, and thus mimics a real 3D convolution kernel. These methods successfully leverage geometric information by introducing a receptive field along the depth-axis. And they can be easily incorporated into pre-trained CNNs while not bring much computational cost. However, in these methods, the depth receptive fields are determined by pre-defined hyperparameters.

In different environments, the scene structures and depth quality can be very different. For example, NYUDv2[27] consists of indoor scenes and captures depth data by Kinect, which has good accuracy. Cityscapes[7] uses a stereo

camera to capture outdoor street scenes, resulting in a much longer depth range and noisier depth data. Naturally, the receptive field along the depth-axis should not be the same across different environment settings. If we artificially pre-define depth receptive fields for different environments, it would bring many parameter-adjusting works, and the selected parameters are possibly still not suitable for the specific environment. Therefore, we need a method that can not only build a receptive field along the depth-axis to handle 3D geometry, but also flexibly learn the receptive field for different environments.

To address the aforementioned problems, in this paper, we propose a novel convolution operator called malleable 2.5D convolution (illustrated in Fig. 1). Similar to 2.5D convolution, the malleable 2.5D convolution can have either one or more 2D convolution kernels sequentially arranged along the depth-axis. To determine the depth receptive field of each kernel, we adopt a softmax classification to assign each pixel to kernels according to pixels’ relative depth differences. The assigning process is differentiable and can be learnt by gradient descent. We also introduce learnable “kernel rebalancing weights” parameters to rebalance the output scale of each kernel, since pixels are not distributed evenly in each class. The malleable 2.5D convolution can flexibly learn the depth receptive field for different environments (as shown in Fig. 4 while only introducing a small number of additional parameters ($2k + 3$ parameters if it has k kernels)). Meanwhile, because malleable 2.5D convolutions are based on 2D convolution kernels, they can be easily incorporated into pre-trained 2D CNNs by simply replacing standard 2D convolutions.

Our contributions can be summarized as follows:

- We propose a novel convolution operator called malleable 2.5D convolution that has learnable receptive fields along the depth-axis.
- Two techniques are proposed in the malleable 2.5D convolution: 1) We propose a differentiable pixel assigning method to achieve learnable depth receptive field; 2) We introduce “kernel rebalancing weights” parameters to rebalance the uneven pixel distribution in the kernels.
- We conduct extensive experiments on both indoor RGB-D semantic segmentation dataset NYUDv2[27] and outdoor dataset Cityscapes[7], and validate the effectiveness and generalization ability of our method.

2 Related Works

RGB-D Scene Parsing Benefiting from the great success of deep convolutional networks[15, 28, 11], Fully Convolutional Networks (FCNs)[21] and its successors[2, 19, 4, 26, 38, 36] have achieved promising results for RGB semantic segmentation. RGB-D segmentation extends RGB semantic segmentation by providing additional depth data. A widely applied method is to encode depth into HHA features[9]: horizontal disparity, height above the ground and angle with gravity direction. It is usually used in two-stream style networks[25, 9, 17, 10, 16, 35] to process RGB images and HHA images and fuse the features or predictions. Other methods attempt to exploit geometric clues from depth data

instead of treating them as features. Some works[29, 39, 23] transform RGB-D images into 3D data and use 3D networks to handle geometry. Other works [18, 5, 14, 13] take advantage of the fact that the scales of objects in images are inversely proportional to the depths, and change the receptive fields of convolutions according to the depth information.

Depth-aware CNN[31] and 2.5D Convolution[34] are two methods most related to our work. Depth-aware CNN applies a mask to construct a soft receptive field along the depth-axis where more distant pixels are weakened and nearer pixels have more contribution to the output. 2.5D convolution seeks to mimic a 3D convolution kernel with several 2D convolution kernels on 2D plane. It builds a grid receptive field in 3D space, and accordingly assigns each pixel to one of the kernels in a similar way a 3D convolution does. Both methods pre-define the receptive field along the depth-axis by hyperparameters, making them rely on parameter adjusting and therefore more difficult to be applied in different environments. Our method solves this problem by building a learnable depth receptive field, and yields better effectiveness and generalization ability.

Convolutions with Learnable Receptive Fields In 2D CNNs, there are several works attempting to break the fixed grid kernel structure of convolution and construct learnable receptive fields. Spatial Transformer Networks[12] warps feature maps with a learned global spatial transformation. SAC[37] learns the scale of receptive fields according to the contents of local features. Deformable convolution[8, 40] learns 2D offsets and adds them to the regular grid sampling locations in the standard convolution. These works change their receptive field on the 2D plane, which can be implemented by linear transformation and interpolation. However, along the depth-axis, we are facing a very different problem. We cannot determine a sample point and then calculate the corresponding feature because pixels in 3D space are very sparse. We have to do it inversely, assigning pixels to sample points. And that makes the process more difficult to be differentiable. Despite the difficulty, our method manages to construct a learnable receptive field along the depth-axis in a novel way. Another difference is that in this work, we aim to learn the depth receptive field suitable for the given dataset instead of each pixel.

Neural Architecture Search Some neural architecture search works[1, 20, 33, 30] try to determine the best convolution kernel shapes (or their combination) for the specific dataset. They use different methods including searching techniques or differentiable optimization. Our method can also be considered as searching the kernel shapes of malleable 2.5D convolutions, but we only search the receptive fields along the depth-axis and we achieve it through learning by gradient descent.

3 Malleable 2.5D Convolution

In this section, we describe and analyze the proposed malleable 2.5D convolution. We firstly introduce the method to construct learnable receptive fields along the depth-axis. Then we introduce the kernel rebalancing mechanism in the operator. Finally, we give some analysis of malleable 2.5D convolution and present how we integrate it into pre-trained CNN.

3.1 Learning Receptive Fields along the Depth-axis

Relative Depth Differences First of all, we define the relative depth differences that are used when we construct receptive fields along the depth-axis. For a convolution kernel, if input map has a downsampling rate r_{down} to the original input image, and the dilation rate is r_{dilate} , then the distance between two neighborhood points sampled by the convolution is $\Delta d_p = r_{down} * r_{dilate}$. Note that Δd_p is a distance on the original input image plane.

For a pixel p_i whose coordinate at the input image is \mathbf{c}_i , and the depth at p_i is $\mathbf{d}(\mathbf{c}_i)$, the coordinate of p_i in 3D space is $\hat{\mathbf{c}}_i = (\mathbf{c}_i - (c_x, c_y)) * \mathbf{d}(\mathbf{c}_i)/f$, where c_x, c_y are the coordinates of the principal point and f is the focal length. c_x, c_y, f are all given by camera parameters (usually two focal lengths f_x, f_y are given, but in most cases they are very close and we can approximately treat them as the same value).

If we project the local sampling grid of the convolution kernel around p_i to 3D space, then the distance between two neighborhood points would be $\Delta d_s(\mathbf{c}_i) = \Delta d_p * \mathbf{d}(\mathbf{c}_i)/f$. The grid size along the depth-axis should be the same with the other two directions. Therefore, we define $\Delta d_s(\mathbf{c}_i)$ as the unit of relative depth difference at image coordinate \mathbf{c}_i . And the relative depth difference between pixel at \mathbf{c}_j and \mathbf{c}_i within the local grid centered at \mathbf{c}_i is given by

$$d(\mathbf{c}_i, \mathbf{c}_j) = (\mathbf{d}(\mathbf{c}_i) - \mathbf{d}(\mathbf{c}_j)) / \Delta d_s(\mathbf{c}_i) \quad (1)$$

Depth Receptive Field Functions For each pixel p_i whose coordinate at the input image is \mathbf{c}_i , a standard 2D convolution performs a weighted sum within the receptive field around p_i :

$$\mathbf{y}(\mathbf{c}_i) = \sum_{\mathbf{c}_p \in \mathcal{R}_p} \mathbf{w}(\mathbf{c}_p) \cdot \mathbf{x}(\mathbf{c}_i + \mathbf{c}_p), \quad (2)$$

where \mathcal{R}_p is the local grid that describes the 2D receptive field around p_i in the input \mathbf{x} , and \mathbf{w} is the convolution kernel. Typically, \mathcal{R}_p is a regular grid defined by kernel size and dilation rate.

A malleable 2.5D convolution has K convolution kernels whose receptive fields are sequentially arranged along the depth-axis. And the pixels are assigned to the kernels according to depth map \mathbf{d} :

$$\mathbf{y}(\mathbf{c}_i) = \sum_{k=1}^K \sum_{\mathbf{c}_p \in \mathcal{R}_p} g_k(\mathbf{d}(\mathbf{c}_i), \mathbf{d}(\mathbf{c}_i + \mathbf{c}_p)) \cdot \mathbf{w}_k(\mathbf{c}_p) \cdot \mathbf{x}(\mathbf{c}_i + \mathbf{c}_p), \quad (3)$$




Fig. 2. Image of h_k and g_k when $K = 3$, $[a_0, a_1, a_2, a_3, a_4] = [-2, -1, 0, 1, 2]$ and $t = 1$. Best view in color

where g_k is the assigning function for kernel k . g_k defines the depth receptive field of kernel k , and satisfies

$$\sum_{k=1}^K g_k(\mathbf{d}(\mathbf{c}_i), \mathbf{d}(\mathbf{c}_i + \mathbf{c}_p)) \leq 1, \forall \mathbf{c}_p \in \mathcal{R}_p. \quad (4)$$

To make the assigning functions differentiable, we implement them as a softmax classification:

$$g_k(\mathbf{d}(\mathbf{c}_i), \mathbf{d}(\mathbf{c}_i + \mathbf{c}_p)) = \frac{\exp(h_k(\mathbf{d}(\mathbf{c}_i), \mathbf{d}(\mathbf{c}_i + \mathbf{c}_p)))}{\sum_{m=0}^{K+1} \exp(h_m(\mathbf{d}(\mathbf{c}_i), \mathbf{d}(\mathbf{c}_i + \mathbf{c}_p)))}. \quad (5)$$

Here h_0 and h_{K+1} are the functions for the two classes outside of all receptive fields (in front of and behind). For $k = 1, 2, \dots, K$, h_k are defined by the relative depth difference as

$$h_k(\mathbf{d}(\mathbf{c}_i), \mathbf{d}(\mathbf{c}_i + \mathbf{c}_p)) = -(d(\mathbf{c}_i, \mathbf{c}_i + \mathbf{c}_p) - a_k)^2/t, \quad (6)$$

where a_k is a learnable parameter that determines the center of the kernel's depth receptive field, and t is a learnable temperature parameter that can sharpen/soften the activation of softmax. h_0 and h_{K+1} are defined as

$$\begin{aligned} h_0 &= -\text{sgn}(d(\mathbf{c}_i, \mathbf{c}_i + \mathbf{c}_p) - a_0) \cdot (d(\mathbf{c}_i, \mathbf{c}_i + \mathbf{c}_p) - a_0)^2/t, \\ h_{K+1} &= \text{sgn}(d(\mathbf{c}_i, \mathbf{c}_i + \mathbf{c}_p) - a_{K+1}) \cdot (d(\mathbf{c}_i, \mathbf{c}_i + \mathbf{c}_p) - a_{K+1})^2/t. \end{aligned} \quad (7)$$

Here sgn denotes the signum function, which is used to make h_0 and h_{K+1} monotonic so that they construct borders of the receptive field.

The constructed depth receptive fields are controlled by parameters t and $a_k, k = 0, 1, \dots, K+1$. And both sets of parameters are learnable by gradient descent. Fig 2 gives an example of the image of h_k and g_k .




Fig. 3. The ratio of pixels assigned to each kernel, before and after rebalance. The figure shows the malleable 2.5D convolution at res5 stage of a trained model. We count the sum of g_k and $s_k \cdot g_k$ for each kernel across the whole NYUDv2 dataset and calculate the ratio. More cases will be presented in the supplementary material.

3.2 Kernel Rebalancing

Through the receptive field functions g_k , we can assign pixels to convolution kernels. However, pixels are not evenly distributed along the depth-axis. In other words, the expectation $\mathbb{E}[g_k(\mathbf{d}(\mathbf{c}_i), \mathbf{d}(\mathbf{c}_i + \mathbf{c}_p))]$ are not equal for different parameters of g_k . This implies that the outputs of different kernels might have different value scales, and the scales change with parameters of g_k . Theoretically, the scale change can be adjusted by \mathbf{w}_k , but implicitly learning a scale factor in convolution kernel weights could increase learning difficulty. Therefore, we introduce a scale factor s_k to rebalance the output scales of different kernel. Then Eq. 3 is modified as

$$\mathbf{y}(\mathbf{c}_i) = \sum_{k=1}^K \sum_{\mathbf{c}_p \in \mathcal{R}_p} s_k \cdot g_k(\mathbf{d}(\mathbf{c}_i), \mathbf{d}(\mathbf{c}_i + \mathbf{c}_p)) \cdot \mathbf{w}_k(\mathbf{c}_p) \cdot \mathbf{x}(\mathbf{c}_i + \mathbf{c}_p), \quad (8)$$

And s_k -s must satisfy $s_k \geq 0, \forall k = 1, 2, \dots, K$ and $\sum_{k=1}^K s_k = 1$. Note that $\sum_{k=1}^K s_k$ can actually equal to any constant considering that in modern CNNs, a convolution layer is almost always followed by a normalization layer. To ensure that s_k -s satisfy the conditions, we implement them as

$$s_k = \frac{\exp(b_k)}{\sum_{k=1}^K \exp(b_k)} \quad (9)$$

And the malleable 2.5D convolution learns b_k by gradient descent.

Fig. 3 shows the effect of kernel rebalancing. We can see that the kernel rebalancing mechanism meets our expectation and successfully rebalance the pixel distribution in different kernels.

3.3 Understanding Malleable 2.5D Convolution

Comparisons with other RGB-D convolutions To get a better understanding of learnable depth receptive fields, we compare our malleable 2.5D convolution with the two previous RGB-D convolutions, Depth-aware Convolution and the 2.5D Convolution. These two convolutions can both be written as the form of Eq. 3. In Depth-aware Convolution, K is 1, and g_1 is defined as

$$g_1(\mathbf{d}(\mathbf{c}_i), \mathbf{d}(\mathbf{c}_i + \mathbf{c}_p)) = \exp(-\alpha|\mathbf{d}(\mathbf{c}_i) - \mathbf{d}(\mathbf{c}_i + \mathbf{c}_p)|), \quad (10)$$

where α is a pre-defined constant. And in 2.5D convolution, K can be any positive integer, and g_k is defined as

$$g_k(\mathbf{d}(\mathbf{c}_i), \mathbf{d}(\mathbf{c}_i + \mathbf{c}_p)) = \begin{cases} 1, & k - 1 - \frac{K}{2} \leq d(\mathbf{c}_i, \mathbf{c}_i + \mathbf{c}_p) < k - \frac{K}{2} \\ 0, & \text{otherwise} \end{cases} \quad (11)$$

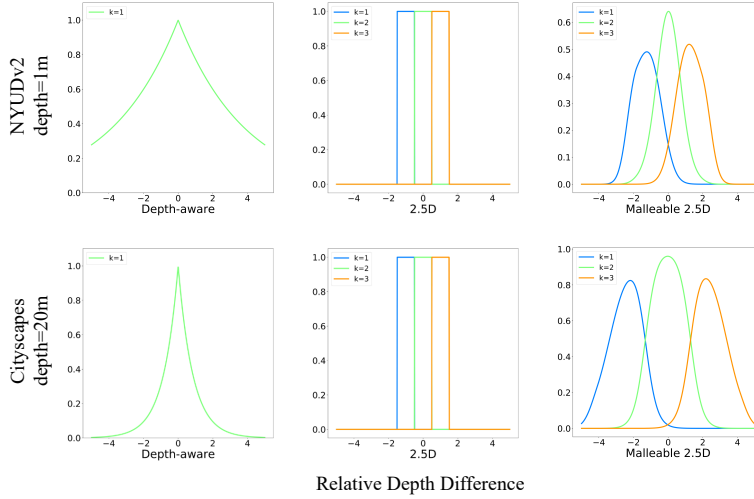


Fig. 4. Comparison of depth receptive field functions g_k . Best view in color. We compare the depth receptive fields at res5 stage after training. The first row shows a case in NYUDv2 where the depth $\mathbf{d}(\mathbf{c}_i) = 1m$. The second row shows a case in Cityscapes where the depth $\mathbf{d}(\mathbf{c}_i) = 20m$. Note that we scale the y-axis to see better. The overall scale does not affect output results because of batch normalizations

In Fig 4, we compare depth receptive fields of these two convolutions and our method. We visualize depth receptive fields in two typical cases in NYUDv2 and Cityscapes respectively. In the two datasets, both focal lengths and the depth ranges are different. The Depth-aware Convolution uses absolute depth difference to control the receptive field. When the depths and focal length changes,

the 3D distance between two adjacent pixels on the image could be very different. And we can see that Depth-aware Convolution cannot fit both cases with the same parameter setting. 2.5D convolution adopts relative depth differences, which makes it more robust to varying depth and focal length. However, 2.5D convolution’s receptive field along the depth-axis is also pre-determined and cannot adapt to different datasets without handcraft adjusting. In contrast, malleable 2.5D convolution not only can avoid the influence of varying depth and focal length, but also automatically learn the depth receptive field for different environments. As an outdoor dataset, Cityscapes has a larger range of depth. And the depth maps in Cityscapes are relatively noisy, not as sharp as those in NYUDv2. Intuitively, the depth receptive field for Cityscapes should be larger and less sharp. We can see from the figure that malleable 2.5D convolution learns a wider depth receptive field for Cityscapes than that for NYUDv2, and the edges drop slower.

Table 1. Comparison of the computational cost of different convolutions. The input size is 768×768 . "kernels" means the kernel number of used RGB-D convolutions

Method	kernels	FLOPs(G)	Params(M)
Baseline	1	215.673	59.468
Depth-aware[31]	1	215.675	59.468
Malleable 2.5D	1	215.679	59.468
2.5D[34]	3	234.701	65.734
Malleable 2.5D	3	234.711	65.734

In Table 1, we compare the FLOPs and parameter number of different convolutions. We present ResNet-based DeepLabv3+[4] as the baseline, and replace the 3×3 convolution with a RGB-D convolution in the first residual unit in each stage of the ResNet. It shows that when using the same number of kernels, the malleable 2.5D convolution only brings very minor additional computational cost to achieve learnable depth receptive fields.

Usage of Malleable 2.5D Convolution Malleable 2.5D convolutions can be easily incorporated into CNNs by replacing standard 2D convolutions. The inputs of malleable 2.5D convolutions are standard feature maps, depth maps, and camera parameters. And the outputs of malleable 2.5D convolutions are the same with standard 2D convolutions.

In RGB-D semantic segmentation, utilizing pre-trained models is essential to attain good performance. When replacing standard convolutions with malleable 2.5D convolutions, we do not abandon the pre-trained weights of the original convolutions but adopt a simple parameter loading strategy to make use of them. We duplicate the pre-trained weights and load them into each kernel of

the malleable 2.5D convolution. in finetuning time, the k kernels start from the same initialization and gradually learn to model geometric relations.

4 Experiments

4.1 Datasets

We evaluate our method on two popular RGB-D scene parsing datasets: NYUDv2[27] and Cityscapes[7]. These two datasets respectively contains indoor and outdoor scenes and the depth sources are different. Therefore, the depth ranges, object sizes and the quality of depth data are very different in them. Evaluating on these two datasets is challenging and can validate the robustness and generalization ability of RGB-D scene parsing methods. We will introduce these two datasets in detail in the supplementary material.

4.2 Implementation Details

Model Implementation DeepLabv3+[4] is a widely recognized state-of-the-art method for semantic segmentation. We adopt ResNet-based DeepLabv3+ pre-trained on ImageNet[24] as our baseline network. And for NYUDv2, we moreover adopt a multi-stage merging block inspired by previous works[19, 16]. Details of the network structures will be illustrated in the supplementary material. To evaluate the our method, we replace the 3×3 convolution with a malleable 2.5D convolution in the first residual unit in each stage of the ResNet.

Our implementation is based on PyTorch[22]. Synchronized batch normalization are adopted for better batch statistics. By default, we use malleable 2.5D convolution with 3 kernels, and the parameter $[a_0, a_1, a_2, a_3, a_4]$ are initialized as $[-2, -1, 0, 1, 2]$, t is initialized as 1, and b_k -s are all initialized as 0.

Training Settings We use SGD optimizer with momentum to train our model. The momentum is fixed as 0.9 and the weight decay is set to 0.0001. We employ a "poly" learning rate policy where the initial learning rate is multiplied by $(1 - \frac{iter}{max_iter})^{power}$. The initial learning rate is set to 0.01 and the power is set to 0.9. We use batch size of 16 and train our model for 40k iterations for NYUDv2 and 60k iterations for Cityscapes. For data augmentation, we use random cropping, random horizontal flipping and random scaling with scale $\in \{0.75, 1, 1.25, 1.5, 1.75, 2\}$. Besides, we adopt the bootstrapped cross-entropy loss as in [32] for Cityscapes experiments.

4.3 Main Results

First of all, we compare our method with previous RGB-D convolutions to validate the benefits of learnable receptive fields along the depth-axis. The results are shown in Table 2. We compare Depth-aware, 2.5D and malleable 2.5D convolution on both NYUDv2 and Cityscapes datasets. When implementing the

Table 2. Comparison with other RGB-D convolutions on NYUDv2 and Cityscapes. The backbone model is ResNet-50. "kernels" means the kernel number of used RGB-D convolutions

Method	kernels	NYUDv2		Cityscapes	
		mIoU(%)	pixel Acc(%)	mIoU(%)	pixel Acc(%)
Baseline	1	44.56	73.01	79.94	96.34
Depth-aware[31]	1	46.69	74.27	79.01	96.32
Malleable 2.5D	1	47.08	75.13	80.26	96.40
2.5D[34]	3	48.23	75.73	78.63	96.29
Malleable 2.5D	3	48.80	76.03	80.81	96.51

Table 3. Comparison with state-of-the-art RGB-D scene parsing methods on NYUDv2. Multi-scale and flipping inference strategies are used when evaluating our method.

Method	Backbone	mIoU(%)	pixel Acc(%)
FCN+HHA[25]	VGG-16 $\times 2$	34.0	65.4
Depth-aware+HHA[31]	VGG-16 $\times 2$	43.9	-
CFN(RefineNet-152)[18]	ResNet-152 $\times 2$	47.7	-
2.5D[34]	ResNet-101 $\times 1$	48.4	75.3
2.5D+HHA[34]	ResNet-101 $\times 2$	49.1	75.9
RDF-101[16]	ResNet-101 $\times 2$	49.1	75.6
RDF-152[16]	ResNet-152 $\times 2$	50.1	76.0
Idempotent[35]	ResNet-101 $\times 2$	50.6	76.3
Malleable 2.5D	ResNet-50 $\times 1$	49.7	76.3
Malleable 2.5D	ResNet-101 $\times 1$	50.9	76.9




Fig. 5. Segmentation results on NYUDv2 test dataset. Black regions in "Ground Truth" are the ignoring category. Models are based on ResNet-50

other two methods, we use the default parameters provided in their papers[31, 34] which are tuned on NYUDv2. All evaluations are conducted without multi-scale and flip testing tricks. For fairness, we compare malleable 2.5D convolution with other convolutions in the case that they have the same kernel number. And the single-kernel version of malleable 2.5D convolution’s parameter $[a_0, a_1, a_2]$ are initialized as $[-1, 0, 1]$. The results show that the malleable 2.5D convolution consistently outperforms other methods and baselines. It is worth noting that Depth-aware and 2.5D convolutions fail on Cityscapes. They decrease the performance even depth information is incorporated. However, our malleable 2.5D convolution still works well on Cityscapes with the same parameter initialization on NYUDv2. It adaptively learns the depth receptive field that fits the outdoor environment and improves segmentation performance. Fig. 5 visualizes qualitative comparison results between the baseline and our method on NYUv2 test set.

NYUDv2 provides high-quality depth data, and most RGB-D scene parsing methods take it as the main benchmark. Therefore, we then compare with more state-of-the-art RGB-D scene parsing methods on NYUDv2, shown in Table 3. As other works do[16, 34, 35], we adopt test-time multi-scale and flip inference to obtain best performance. We evaluate models based on both ResNet-50 and ResNet-101 backbones. And we use malleable 2.5D convolutions with 3 kernels. Different from many works that adopt two-stream network design to process RGB+HHA inputs, we do not use HHA input and thus do not double the computational cost brought by backbones. The results show that even without doubling the backbone, our model achieves promising performance and outperforms other methods.

4.4 Ablation Studies

In this subsection, we conduct ablation studies on NYUDv2 dataset to validate the efficacy of our method.

Number of kernels In Table 4, we conduct experiments that compare different numbers of convolutional kernels in malleable 2.5D convolutions. While the performance improves when the number increases from 1 to 3, the performance of the 5-kernel case drops slightly compared to the 3-kernel case. We suppose the reason is that 3 kernels are enough to model 3D relations between pixels, and in the 5-kernel case there are too few pixels assigned to the two outer kernels and they cannot be sufficiently trained.

Learnable parameters In Table 5, we evaluate the effects of introduced learnable parameters. We fix part of the introduced parameters $\{a_k, t, b_k\}$ at the initialized values, and set rest of them learnable. The results demonstrate that when we fix all the parameters and thus fix the depth receptive fields, the performance is the worst. And the fixed malleable 2.5D convolution gets a very similar result to 2.5D convolution. When only the kernel center a_k is learnable, the

Table 4. Ablation study for number of convolution kernels. The backbone is ResNet-50

Kernel numbers	mIoU(%)	pixel Acc(%)
1	47.08	75.13
3	48.80	76.03
5	48.17	75.59

performance stays almost the same. And when both a_k and t are learnable, the mIoU improves from 48.33% to 48.62% and pixel accuracy improves from 75.74% to 75.91%. We argue that this is because the temperature term T controls the sharpness of depth receptive fields, and therefore it is crucial to set both a_k and t learnable for assigning pixels to different kernels more accurately. When the kernel rebalancing parameter b_k is introduced, the performance moreover improves to the result of complete malleable 2.5D convolution. These results validate that the improvement is indeed brought by the introduced learnable parameters.

Table 5. Ablation study for introduced learnable parameters in malleable 2.5D convolution. The backbone model is ResNet-50

Method	Learnable params	mIoU(%)	pixel Acc(%)
Baseline	-	44.56	73.01
2.5D[34]	-	48.23	75.73
Malleable 2.5D	None	48.28	75.68
	a_k	48.33	75.74
	a_k, t	48.62	75.91
	a_k, t, b_k	48.80	76.03

Initialization of introduced parameters The parameter a_k serves as the center of kernel k 's receptive field along the depth-axis. It is the most direct factor that controls the learnt depth receptive field. In Table 6, we evaluate the effects of different initialization of a_k for 3-kernel malleable 2.5D convolution. When we use large values to initialize a_k , the performance drops slightly. However, both results outperform the baseline and 2.5D convolution. This validates the effectiveness and the robustness of the learnable depth receptive field in our method. More experiments of different parameter initializations will be included in the supplementary material.

Malleable 2.5D Convolution in Different Layers To reveal the effects of using malleable 2.5D convolution in different locations in the backbone network,

Table 6. Results of different initialization of a_k . The backbone model is ResNet-50

Method	Initialization	mIoU(%)	pixel Acc(%)
Baseline	-	44.56	73.01
2.5D[34]	-	48.23	75.73
Malleable 2.5D	$[-4, -2, 0, 2, 4]$	48.66	75.94
	$[-2, -1, 0, 1, 2]$	48.80	76.03

we conduct a series of experiments. The results are shown in Table 7. When using more malleable 2.5D convolution, the network gains more capability to handle 3D geometric information and achieves better segmentation performance.

Table 7. Results of using malleable 2.5D convolutions in different layers. "Replaced Location" means in which stage in the ResNet the first 3×3 convolution is replaced by malleable 2.5D convolution. The backbone model is ResNet-50

Method	Replaced Location	mIoU(%)	pixel Acc(%)
Baseline	-	44.56	73.01
Malleable 2.5D	res2, res3	48.27	75.83
	res4, res5	48.34	75.77
	res3, res4, res5	48.58	75.93
	res2, res3, res4, res5	48.80	76.03

5 Conclusion

We propose a novel RGB-D convolution operator called malleable 2.5D convolution, which has learnable receptive fields along the depth-axis. By learning the receptive field along the depth-axis, malleable 2.5D convolution can learn to adapt to different environments without handcraft parameter adjusting and improve RGB-D scene parsing performance. And malleable 2.5D convolution can be easily incorporated into pre-trained standard CNNs. Our extensive experiments on NYUDv2 and Cityscapes validate the effectiveness and the generalization ability of our method.

Acknowledgments: This work is supported by the National Key Research and Development Program of China (2017YFB1002601, 2016QY02D0304), National Natural Science Foundation of China (61375022, 61403005, 61632003), Beijing Advanced Innovation Center for Intelligent Robots and Systems (2018IRS11), and PEK-SenseTime Joint Laboratory of Machine Vision.

References

1. Chen, L., Collins, M.D., Zhu, Y., Papandreou, G., Zoph, B., Schroff, F., Adam, H., Shlens, J.: Searching for efficient multi-scale architectures for dense image prediction. In: NeurIPS. pp. 8713–8724 (2018)
2. Chen, L., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell.* **40**(4), 834–848 (2018)
3. Chen, L., Papandreou, G., Schroff, F., Adam, H.: Rethinking atrous convolution for semantic image segmentation. *CoRR* **abs/1706.05587** (2017)
4. Chen, L., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: ECCV (7). *Lecture Notes in Computer Science*, vol. 11211, pp. 833–851. Springer (2018)
5. Chen, Y., Mensink, T., Gavves, E.: 3d neighborhood convolution: Learning depth-aware features for RGB-D and RGB semantic segmentation. In: 3DV. pp. 173–182. IEEE (2019)
6. Cheng, Y., Cai, R., Li, Z., Zhao, X., Huang, K.: Locality-sensitive deconvolution networks with gated fusion for RGB-D indoor semantic segmentation. In: CVPR. pp. 1475–1483. IEEE Computer Society (2017)
7. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27–30, 2016. pp. 3213–3223. IEEE Computer Society (2016). <https://doi.org/10.1109/CVPR.2016.350>, <https://doi.org/10.1109/CVPR.2016.350>
8. Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., Wei, Y.: Deformable convolutional networks. In: ICCV. pp. 764–773. IEEE Computer Society (2017)
9. Gupta, S., Girshick, R.B., Arbeláez, P.A., Malik, J.: Learning rich features from RGB-D images for object detection and segmentation. In: ECCV (7). *Lecture Notes in Computer Science*, vol. 8695, pp. 345–360. Springer (2014)
10. Hazirbas, C., Ma, L., Domokos, C., Cremers, D.: FuseNet: incorporating depth into semantic segmentation via fusion-based CNN architecture. In: ACCV (1). *Lecture Notes in Computer Science*, vol. 10111, pp. 213–228. Springer (2016)
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778. IEEE Computer Society (2016)
12. Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks. In: NIPS. pp. 2017–2025 (2015)
13. Kang, B., Lee, Y., Nguyen, T.Q.: Depth-adaptive deep neural network for semantic segmentation. *IEEE Trans. Multimedia* **20**(9), 2478–2490 (2018)
14. Kong, S., Fowlkes, C.C.: Recurrent scene parsing with perspective understanding in the loop. In: CVPR. pp. 956–965. IEEE Computer Society (2018)
15. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS. pp. 1106–1114 (2012)
16. Lee, S., Park, S., Hong, K.: RDFNet: RGB-D multi-level residual feature fusion for indoor semantic segmentation. In: ICCV. pp. 4990–4999. IEEE Computer Society (2017)
17. Li, Z., Gan, Y., Liang, X., Yu, Y., Cheng, H., Lin, L.: LSTM-CF: unifying context modeling and fusion with lstms for RGB-D scene labeling. In: ECCV (2). *Lecture Notes in Computer Science*, vol. 9906, pp. 541–557. Springer (2016)

18. Lin, D., Chen, G., Cohen-Or, D., Heng, P., Huang, H.: Cascaded feature network for semantic segmentation of RGB-D images. In: ICCV. pp. 1320–1328. IEEE Computer Society (2017)
19. Lin, G., Milan, A., Shen, C., Reid, I.D.: RefineNet: multi-path refinement networks for high-resolution semantic segmentation. In: CVPR. pp. 5168–5177. IEEE Computer Society (2017)
20. Liu, H., Simonyan, K., Yang, Y.: DARTS: differentiable architecture search. In: ICLR (Poster). OpenReview.net (2019)
21. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR. pp. 3431–3440. IEEE Computer Society (2015)
22. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in PyTorch. In: NIPS Autodiff Workshop (2017)
23. Qi, X., Liao, R., Jia, J., Fidler, S., Urtasun, R.: 3d graph neural networks for RGBD semantic segmentation. In: ICCV. pp. 5209–5218. IEEE Computer Society (2017)
24. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M.S., Berg, A.C., Li, F.: Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* **115**(3), 211–252 (2015)
25. Shelhamer, E., Long, J., Darrell, T.: Fully convolutional networks for semantic segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**(4), 640–651 (2017)
26. Shen, F., Gan, R., Yan, S., Zeng, G.: Semantic segmentation via structured patch prediction, context CRF and guidance CRF. In: CVPR. pp. 5178–5186. IEEE Computer Society (2017)
27. Silberman, N., Hoiem, D., Kohli, P., Fergus, R.: Indoor segmentation and support inference from RGBD images. In: ECCV (5). *Lecture Notes in Computer Science*, vol. 7576, pp. 746–760. Springer (2012)
28. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *CoRR* **abs/1409.1556** (2014)
29. Song, S., Yu, F., Zeng, A., Chang, A.X., Savva, M., Funkhouser, T.A.: Semantic scene completion from a single depth image. In: CVPR. pp. 190–198. IEEE Computer Society (2017)
30. Tan, M., Le, Q.V.: Mixconv: Mixed depthwise convolutional kernels. *CoRR* **abs/1907.09595** (2019)
31. Wang, W., Neumann, U.: Depth-aware CNN for RGB-D segmentation. In: ECCV (11). *Lecture Notes in Computer Science*, vol. 11215, pp. 144–161. Springer (2018)
32. Wu, Z., Shen, C., van den Hengel, A.: High-performance semantic segmentation using very deep fully convolutional networks. *CoRR* **abs/1604.04339** (2016)
33. Xie, S., Zheng, H., Liu, C., Lin, L.: SNAS: stochastic neural architecture search. In: ICLR (Poster). OpenReview.net (2019)
34. Xing, Y., Wang, J., Chen, X., Zeng, G.: 2.5 d convolution for rgb-d semantic segmentation. In: 2019 IEEE International Conference on Image Processing (ICIP). pp. 1410–1414. IEEE (2019)
35. Xing, Y., Wang, J., Chen, X., Zeng, G.: Coupling two-stream rgb-d semantic segmentation network by idempotent mappings. In: 2019 IEEE International Conference on Image Processing (ICIP). pp. 1850–1854. IEEE (2019)
36. Yu, F., Koltun, V.: Multi-scale context aggregation by dilated convolutions. *CoRR* **abs/1511.07122** (2015)

37. Zhang, R., Tang, S., Zhang, Y., Li, J., Yan, S.: Scale-adaptive convolutions for scene parsing. In: IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017. pp. 2050–2058. IEEE Computer Society (2017). <https://doi.org/10.1109/ICCV.2017.224>, <https://doi.org/10.1109/ICCV.2017.224>
38. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: CVPR. pp. 6230–6239. IEEE Computer Society (2017)
39. Zhong, Y., Dai, Y., Li, H.: 3d geometry-aware semantic labeling of outdoor street scenes. In: ICPR. pp. 2343–2349. IEEE Computer Society (2018)
40. Zhu, X., Hu, H., Lin, S., Dai, J.: Deformable convnets V2: more deformable, better results. In: CVPR. pp. 9308–9316. Computer Vision Foundation / IEEE (2019)