

# Dilated Convolutions with Lateral Inhibitions for Semantic Image Segmentation

Yujiang Wang<sup>1</sup> Mingzhi Dong<sup>2</sup> Jie Shen<sup>1,3,†</sup> Yiming Lin<sup>1,3</sup> Maja Pantic<sup>1,3</sup>  
<sup>1</sup>Imperial College London <sup>2</sup>University College London <sup>3</sup>Facebook London

yujiang.wang14@imperial.ac.uk mingzhi.dong.13@ucl.ac.uk  
{jie.shen07, yiming.lin15, m.pantic}@imperial.ac.uk

## Abstract

*Dilated convolutions are widely used in deep semantic segmentation models as they can enlarge the filters' receptive field without adding additional weights nor sacrificing spatial resolution. However, as dilated convolutional filters do not possess positional knowledge about the pixels on semantically meaningful contours, they could lead to ambiguous predictions on object boundaries. In addition, although dilating the filter can expand its receptive field, the total number of sampled pixels remains unchanged, which usually comprises a small fraction of the receptive field's total area. Inspired by the Lateral Inhibition (LI) mechanisms in human visual systems, we propose the dilated convolution with lateral inhibitions (LI-Convs) to overcome these limitations. Introducing LI mechanisms improves the convolutional filter's sensitivity to semantic object boundaries. Moreover, since LI-Convs also implicitly take the pixels from the laterally inhibited zones into consideration, they can also extract features at a denser scale. By integrating LI-Convs into the Deeplabv3+ architecture, we propose the Lateral Inhibited Atrous Spatial Pyramid Pooling (LI-ASPP), the Lateral Inhibited MobileNet-V2 (LI-MNV2) and the Lateral Inhibited ResNet (LI-ResNet). Experimental results on three benchmark datasets (PASCAL VOC 2012, CelebAMask-HQ and ADE20K) show that our LI-based segmentation models outperform the baseline on all of them, thus verify the effectiveness and generality of the proposed LI-Convs.*

## 1. Introduction

Since the introduction of the pioneering Fully Convolutional Networks (FCN) [28], deep Convolutional Neural Networks (CNNs) [9, 57, 27, 56, 20] have made impressive progress in semantic image segmentation, a task that performs per-pixel classifications. In deep CNN models, a series of convolutions and spatial poolings are applied to obtain progressively more abstract and more representative

feature descriptors with decreasing resolutions. As a consequence, the deepest features can have significantly lower resolution than the original image (e.g. only 1/16 or 1/32 of the input size in FCN [28]), hence it would be difficult to decode these features into the segmentation map at the same size of the input image without losing details. This is a crucial challenge in the semantic segmentation task.

Dilated convolutions [19], which are first applied to the semantic segmentation task by [54, 6], can effectively overcome such difficulties and thus are widely employed in state-of-the-art segmentation methods [27, 9, 53, 5, 48]. By inserting zeros (*dilation*) into the convolutional filters, dilated convolutions can observe features from larger areas without increasing the kernel parameters, which is important to the extractions of global semantic features. Besides, it can also produce feature maps that are invariant input resolutions. In practice, dilated convolutions can be utilised to retain the resolution of the feature maps when encoding representations in the backbone network [55, 48], typically by replacing certain convolutional layers with dilated ones. It can also be employed during the decoding stage to generate more robust semantic labels, e.g. the Atrous Spatial Pyramid Pooling (ASPP) [8, 7] adopts three parallel dilated convolutions with different dilation rates to aggregate the multi-scale contextual information.

Despite its broad applications, dilated convolutions still have several limitations. The pixels around semantically meaningful contours separate different objects and possess stronger semantic information. In dilated convolution, however, the importance of those pixels are not explicitly accentuated, and therefore such positional significance has to be implicitly learnt. This can lead to ambiguous and misleading boundary labels. Various approaches have been proposed to compensate for such problems and to refine the contour predictions, including the Conditional Random Fields (CRF) [7, 4] and the decoder component in Deeplabv3+ [9]. However, dilated convolution's sensitivity on spotting semantically meaningful edges still leaves room for improvement.

Additionally, although the receptive field of dilated filters is enlarged, the total number of sampled pixels stay the

<sup>†</sup>Corresponding author.

same, which only consist of a small fraction of pixels in the area. The sparse sampling can somehow impair the potentials for dense prediction tasks like semantic segmentation. Similar concerns were address in [43, 48, 11, 52], and the proposed improvements include a denser Gaussian sampling process [43], a hybrid dilated convolution module [48] and the deformable convolutional filters [11].

In this paper, we propose to overcome the drawbacks in dilated convolutions from a biologically-inspired perspective, which is to leverage the Lateral Inhibition (LI) mechanism in the human visual system. Lateral inhibition [17, 38, 47] is a neurobiological phenomenon that a neuron’s excitation to a stimulus can be suppressed by the activation of its surrounding neurons. Because of the LI mechanism, our retina cells are sensitive to the spatially varying stimulus such as the semantic borderlines between objects, which is crucial to the inborn segmentation abilities of our eyes. See Fig. 1 (Left) for an intuitive illustration of the LI mechanism.

Motivated by such observations, we propose a dilated convolution with lateral inhibitions (LI-Convs) to enhance the convolutional filter’s sensitivity to semantic contours. The LI-Convs also sample the receptive window in a denser fashion by implicitly making inferences on pixels within the lateral inhibited zones. To evaluate LI-Convs, we follow the Deeplabv3+ [9] segmentation models and present three LI-based variants which are 1). the Lateral Inhibited Atrous Spatial Pyramid Pooling (LI-ASPP) for decoding semantic features, 2). the Lateral Inhibited MobileNet-V2 (LI-MNV2) and 3). the Lateral Inhibited ResNet (LI-ResNet) as the backbone networks for encoding features. The performance of LI-ASPP, LI-MNV2 and LI-ResNet surpasses the baseline on three segmentation benchmark dataset: PASCAL VOC 2012 [12], CelebAMask-HQ [24] and ADE20K [58], which verifies the effectiveness and generality of the proposed LI-Convs.

## 2. Related Works

**Semantic Image Segmentation** Fully Convolutional Networks (FCN) [28] is the pioneering work of using deep models for semantic segmentation. The fully connected layers in deep image classification models are replaced with convolutional ones to produce semantic heat maps for segmentation predictions. The resolution of such heat maps is typically much smaller than that of the input image (e.g.  $1/32$ ), and various works are proposed to compensate the information loss during decoding such features, including the de-convolutional layers [35, 39, 36], the skip-connections of low-level features [2, 16] and dilated convolutions [54, 8, 53, 27, 5]. Yu *et al.* [54] stacks dilated convolutional layers with different dilation rates in a cascaded manner, leading to a context module for aggregating the multi-scale contextual information. Deeplabv3 [8] builds an Atrous Spatial Pyramid Pooling (ASPP) module consisting

of three parallel dilated convolutions, one  $1*1$  convolution and one image-level pooling, and it also employs dilated convolutions in the backbone network. DenseASPP [53] introduces dense connection into the ASPP module to enlarge its receptive fields and to acquire denser feature pyramid, while the technique of Neural Architecture Search [60] is utilised by [5] to search for an optimal decoding structure of organising dilated convolutions layers. For other segmentation practice [37, 50, 29, 49], readers are referred to [33] for more details.

**Dilated Convolutions** Dilated convolutions, also known as atrous convolutions, is first introduced by Holschneider *et al.* [19] in signal analysis and have broad applications such as object detection [25, 34], lip-reading [51, 32, 30] and optical flow [59, 46]. It is first applied to semantic segmentation by authors of [54, 6] to enlarge filter’s receptive fields without sacrificing the spatial resolution. Conditional Random Fields (CRF) are involved in [4, 7] as a post-processing procedure to refine the ambiguous semantic contour predictions. Similar ideas can be found in Deeplabv3+ [9], which designs a decoding module to incorporate low-level backbone features to improve the qualities of contouring pixels. Deformable convolutions [11] introduce the *offsets* into the sampling grids of filters to better model the spatial relationships. Gaussian kernels are adopted by [43] to obtain pixels at a wider range in dilated convolutions. Wang *et al.* [48] observe the *gridding* effects brought by the fixed sampling locations in dilated kernels and demonstrate a hybrid dilated convolution with different dilated rates. Different from those approaches, we employ the lateral inhibition (LI) mechanisms [17] to enhance the dilated convolutions’ sensitivity on semantically meaningful contours and to implicitly sample features in a denser fashion.

**Lateral Inhibitions** The study on the eyes of horseshoe crab (*Limulus*) performed by Hartline *et al.* [17] reveals the lateral inhibition (LI) effects in visual systems, where the excitation of neighbouring neurons can suppress a cell’s response to the stimuli. Although lateral inhibitions are mainly studied in the field of neuroscience [40, 45, 38], the computer vision community has also shown interests in this mechanism. The recurrent neural network with lateral inhibitions is studied in [31] and it is shown that LI can improve the robustness and efficiency of the network. Authors of [13] introduce LI into a shallow CNN to improve image classification. Similar ideas can be found in the work for colour video segmentation [14]. Those network architectures are somehow too shallow to be useful for recent methods using deep backbones like MobileNet-V2 (MNV2) [42] or ResNet [18]. The idea of LI can also be found in the Local Response Normalisation (LRN) proposed by AlexNet [23], yet the inhibitions in LRN come from different channels on the same spatial locations, which might not be suitable for

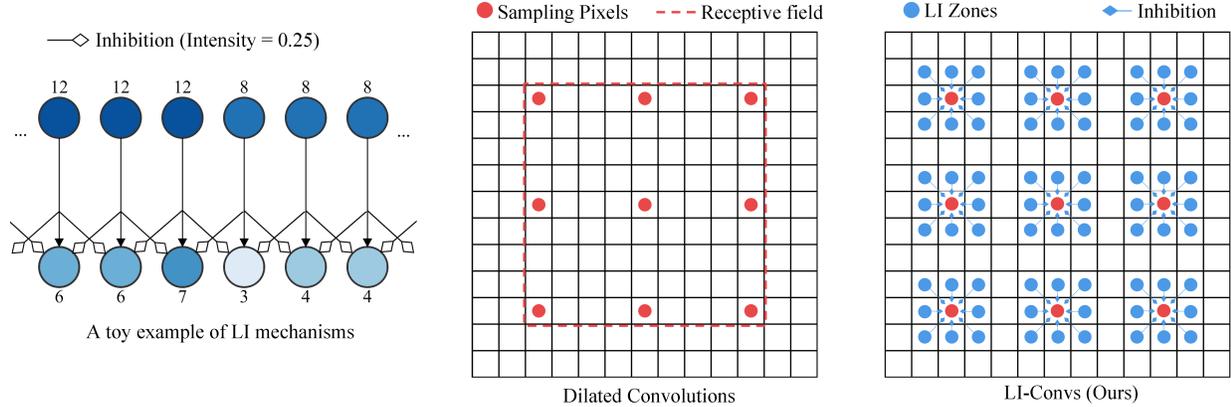


Figure 1. *Left:* A toy example to illustrate the lateral inhibition mechanisms where the LI intensity is set to 0.25. The difference between the two neurons at the centre (representing a semantic contour) becomes more significant after LI. *Middle:* A  $3 \times 3$  convolutional filter where  $d = 4$ . The sampled pixels (denoted as red dots) only comprises a small fraction of all pixels in the receptive field. *Right:* An illustration of the proposed LI-Convs with  $3 \times 3$  lateral inhibition zones. Each sampled pixel receives inhibition signals from eight neighbours to enhance sensitivity on semantic contours and to extract information at a denser scale.

segmentation tasks, also there are no learnable parameters in it. Recently, authors of [3] employ LI in VGG model [44] to improve the performance on saliency detection. However, none of the previous works has evaluated LI’s potentials for semantic segmentation, while their methods of integrating LI do not touch the core mechanisms in deep CNNs such as the convolutional operations. In this work, however, lateral inhibitions work closely with the convolutional filters to fundamentally augment the model’s segmentation powers.

### 3. Dilated Convolutions with Lateral Inhibitions

#### 3.1. Definition

Define  $\Psi_k = \mathbb{Z}^2 \cap [-k, k]^2$  where  $k \in \mathbb{Z}_{\geq 0}$ , and let a discrete function  $F : \Psi_k \mapsto \mathbb{R}$  represents a convolutional filter of size  $(2k + 1)^2$ . Define another discrete function  $G : \mathbb{Z}^2 \mapsto \mathbb{R}$  representing features of arbitrary sizes. Let  $d$  be the dilation rate, a dilated convolutional operator  $\star_d$  is written as

$$(F \star_d G)(\mathbf{p}) = \sum_{d\mathbf{m}+\mathbf{n}=\mathbf{p}} F(\mathbf{m})G(\mathbf{n}) \quad (1)$$

where  $\mathbf{p}, \mathbf{m}, \mathbf{n} \in \mathbb{Z}^2$ . Note that  $\star_d$  turns into a regular convolutional operator when  $d = 1$ , i.e. no dilation is inserted.

With the introduction of lateral inhibitions (LI), the activation of each sampled pixel, i.e.  $G(\mathbf{n})$  in Eq. 1, would be suppressed by its neighbours within a certain range. Let the lateral inhibitions come from a square region of size  $(2t + 1)^2$  centred on  $\mathbf{n}$  where  $t \in \mathbb{Z}_{\geq 0}$ , and refer this region as the lateral inhibition zone (the LI zone). Define  $\Psi_t = \mathbb{Z}^2 \cap [-t, t]^2$  and let  $L : \Psi_t, \mathbb{Z}^2 \mapsto \mathbb{R}$  be a discrete function describing the spatially-varying inhibition intensities in the LI zones, the amount of the total inhibitions received by a sampled pixel  $G(\mathbf{n})$  can be described

as  $\sum_{\mathbf{u}+\mathbf{v}=\mathbf{n}} L(\mathbf{u}, \mathbf{n})G(\mathbf{v})$  where  $\mathbf{u}, \mathbf{v} \in \mathbb{Z}^2$ . Consequently, a dilated convolutional operator with lateral inhibition  $\star_d$  (LI-Convs) can be defined as

$$(F \star_d G)(\mathbf{p}) = \sum_{d\mathbf{m}+\mathbf{n}=\mathbf{p}} F(\mathbf{m})\phi(G(\mathbf{n}) - \sum_{\mathbf{u}+\mathbf{v}=\mathbf{n}} L(\mathbf{u}, \mathbf{n})G(\mathbf{v})) \quad (2)$$

where  $\phi$  represents an activation function like ReLu. The introduced LI terms and non-linearity distinguish LI-Convs in Eq. 2 with Eq. 1. An intuitive comparison between dilated convolutions and the proposed LI-Convs is shown in Fig. 1 (*Middle & Right*).

We can also “dilate” the lateral inhibition zone to efficiently expand its field-of-views, in a similar way to that of dilated convolutions. Consequently, a generalised LI-Convs operator  $\star_d^e$  is defined as

$$(F \star_d^e G)(\mathbf{p}) = \sum_{d\mathbf{m}+\mathbf{n}=\mathbf{p}} F(\mathbf{m})\phi(G(\mathbf{n}) - \sum_{\mathbf{e}\mathbf{u}+\mathbf{v}=\mathbf{n}} L(\mathbf{u}, \mathbf{n})G(\mathbf{v})) \quad (3)$$

where  $e$  denotes the dilation rate in LI zones.

Although a wide variety of kernel forms can be taken by the LI intensity descriptor  $L$ , we opt for an intuitive formulation that is also easy to implement. In particular,  $L(\mathbf{u}, \mathbf{n})$  in Eq. 3 simply takes the production of a differentiable weight  $W_L \in [0, 1] \cap \mathbb{R}$  and an exponentially decaying factor that is related to the distance between  $\mathbf{u}$  and  $\mathbf{n}$ , which can be described as

$$L(\mathbf{u}, \mathbf{n}) = W_L \exp\left(\frac{-D^2(\mathbf{u}, \mathbf{n})}{2\sigma^2}\right) \quad (4)$$

where  $\sigma$  is a parameter representing the standard deviation,  $\exp$  denotes the exponential function and  $D(\mathbf{u}, \mathbf{n})$  refers to a certain distance measurement between  $\mathbf{u}$  and  $\mathbf{n}$ . Here we employ the Euclidean distance.

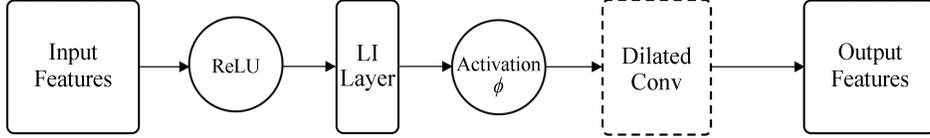


Figure 2. The structure of LI-Convs. The lateral inhibitions is first calculated by the LI layer, and the inhibited features are fed into the dilated convolution layer. The dilated convolution part can be any kind of convolution implementations such as the depthwise one [10].

### 3.2. Implementation of LI-Convs

We take a straight-forward approach to implement the LI-Convs in Eq. 3. We first design a Lateral Inhibition layer (the LI layer) to perform pixel-wise lateral inhibitions, while a dilated convolutional layer is subsequently applied to the inhibited features. The LI layer is essentially a light-weight module that can be flexibly inserted into deep models, while it can be easily implemented as a dilated convolutional layer with specifically shaped filters. In particular, let a discrete function  $K : \Psi_t \mapsto \mathbb{R}$  represent one such LI filter,  $K$  can be described as:

$$K(\mathbf{u}) = \begin{cases} 1.0 & \mathbf{u} = \mathbf{0}. \\ -W_L \exp\left(\frac{-D^2(\mathbf{u}, \mathbf{0})}{2\sigma^2}\right) & \mathbf{u} \neq \mathbf{0}. \end{cases} \quad (5)$$

Note that the LI filter  $K$  has identical size with the LI zones which is  $(2t + 1)^2$ , and applying  $K$  with a stride of 1 can generate pixel-wise inhibited features. We empirically set  $\sigma$  in Eq. 5 to a fixed value during training, thus there is only one weight  $W_L$  to learn for each LI filter, which is significantly less than that of regular convolutional filters. In practice, we learn the lateral inhibition weights in a channel-wise manner, i.e. each LI filter learns a separate  $W_L$ . Therefore, a LI layer will introduce a total of  $C$  learnable weights where  $C$  is the channel number of the input tensor.

A detailed illustration for the LI-Convs implementations can be found in Fig. 2. A ReLU activation is first applied to remove negative neuron response. Then a LI layer with filters in Eq. 5 is employed to extract inhibited features, followed by the activation function  $\phi$  in Eq. 3. A dilated convolution layer of arbitrary form such as the depthwise convolution [10] is subsequently employed.

### 3.3. LI-ASPP, LI-MNV2 and LI-ResNet

We introduce the proposed LI-Convs into the state-of-the-art segmentation model Deeplabv3+ [9] to evaluate the proposed LI-Convs. As shown in Fig. 3, we replace the three  $3 \times 3$  parallel dilated convolution operations in Atrous Spatial Pyramid Pooling (ASPP) [9] with the proposed LI-Convs, leading to the LI-ASPP model. Besides, we also investigate the potentials of LI layer in the backbone network such as the MobileNet-V2 (MNv2) [42] and ResNet [18].

As illustrated in Fig. 4 (Left), we insert the LI layer into the *residual bottleneck* (RB) of MobileNet-V2 [42], which is between the  $1 \times 1$  expansion convolution and  $3 \times 3$  depthwise

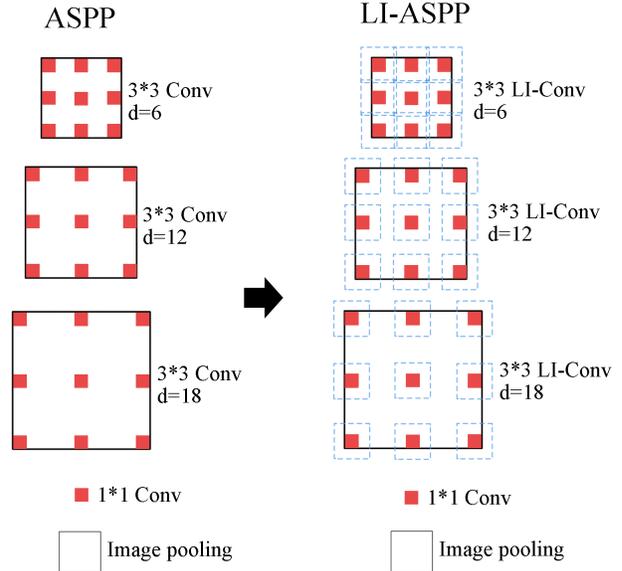


Figure 3. The structures of ASPP and LI-ASPP. ASPP consists of five parallel branches including three dilated convolutions, which are replaced with the proposed LI-Convs in LI-ASPP.

convolution, and we refer the resulting structure as the LI bottleneck layer. In the original MNV2 architecture, there are a total of 17 *residual bottleneck* layers, and we replace the 10<sup>th</sup>, 13<sup>th</sup> and 16<sup>th</sup> (16<sup>th</sup> refers to the second-highest RB layer) *residual bottlenecks* with the LI bottlenecks to obtain the LI-MNV2 network.

Similarly, we modify the *bottleneck unit* (we adopt the one with 3 convolutional layers) in ResNet by inserting a LI layer between the first two weighted layers, as shown in Fig. 4 (Right), and name the new architecture as the LI bottleneck unit. Among those ResNet variants, we select the ResNet-50 architecture in this work and replace its “conv5\_3” layer with the LI bottleneck unit to get the LI-ResNet-50 network.

## 4. Experiments

### 4.1. Datasets

We conduct our experiments on three public benchmark segmentation datasets, which are PASCAL VOC 2012 [12], CelebAMask-HQ [24] and ADE20K [58]. There are a total of 21 semantic classes in PASCAL VOC 2012 dataset [12] which contains 1,464/1,449/1,456 pixel-wise annotated images for train/validation/test. Following [15, 9], we use an

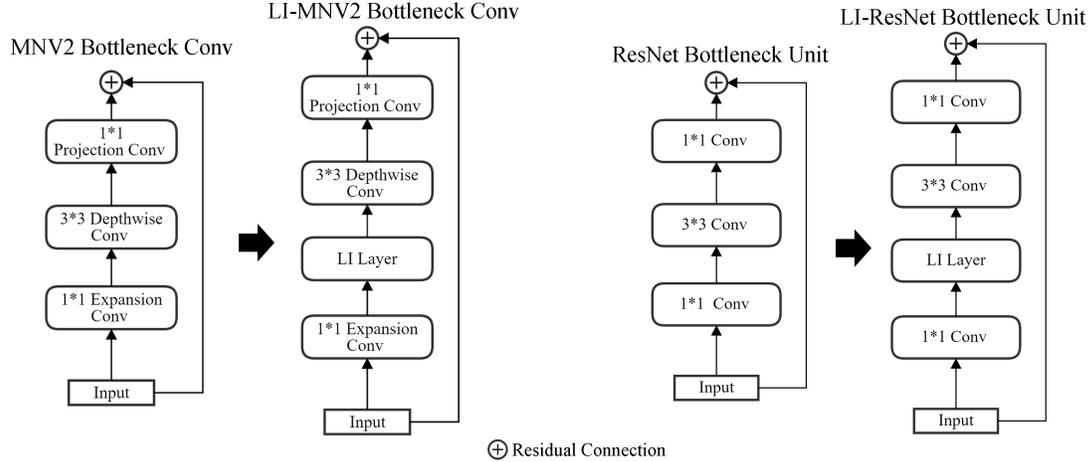


Figure 4. *Left*: The structures of the *residual bottleneck* convolution in MobileNet-V2 and the LI bottleneck. The LI layer is inserted between the  $1 \times 1$  expansion convolution and  $3 \times 3$  depthwise convolution. *Right*: The architecture of the 3-layer *bottleneck unit* in ResNet and the LI bottleneck unit. The LI layer is inserted between the  $1 \times 1$  conv and  $3 \times 3$  conv layers.

augmented train set with a total of 10,582 annotated images. CelebAMask-HQ [24] is a large-scale face parsing dataset with 30,000 pixel-wise labelled face images of 19 classes, and they are split into sets with 24,183/2,993/2,824 images for train, validation and test. ADE20K [58] is a benchmark dataset for scene parsing with 20,210/2,000/3,000 pixel-wise labelled images for train/validation/test. It is a quite challenging dataset, as there are a total of 151 classes in this dataset, and the huge variations of image resolutions also increase the difficulties. We utilise the validation set to evaluate performance on PASCAL VOC 2012 and ADE20K datasets, considering that their test sets are not publicly available, while we follow the standard protocol on CelebAMask-HQ dataset and use the test set for evaluation.

## 4.2. Experimental Setup

**Evaluation metric** Mean Intersection-over-Union (mIoU) is the most widely used evaluation metric for the segmentation task, and we adopt it to evaluate the quality of model predictions. We also report the model parameters and the FLOPs to provide more comprehensive analyses.

**Training Settings** We generally follow the training settings in Deeplabv3+ [9], while we have also made some modifications to suit our needs. Particularly, we use the ImageNet [41] checkpoints provided by the authors of MobileNet-V2 [42] and ResNet [18] to initialise LI-MNV2 and LI-ResNet-50, respectively, while the weights of LI-ASPP are randomly initialised. Note that we do not use the MS COCO dataset [26] to pre-train the model. During training, we set the image crop size to be  $513 \times 513$  for all three datasets, except that we use  $257 \times 257$  crop size when evaluating on ADE20K with the MNV2-based backbone. We train for 120 epochs using a batch size of 16 and Adam [22] is applied to optimise the pixel-wise cross-entropy loss with L2-regularisation. The initial learning rate and the epsilon value in Adam optimiser

are set to 0.0003 and 0.01, respectively. The *output stride*, which is defined in [8] denoting the ratio of original input resolution to the final feature’s resolution, is set to be 16 for all datasets. We adopt strategies in [9, 8] to use the Batch-Norm layers [21] and to randomly scale the training data for augmentation. Depthwise convolution [10] is used in the ASPP implementations following [9]. During evaluations, we set the *output stride* to be 16 for all datasets and employ a single evaluation scale of 1.0, and all crop sizes are set to  $513 \times 513$  except that  $257 \times 257$  is utilised for evaluating MNV2-based backbones on ADE20K.

**LI Layer Settings** A lateral inhibition layer has several key hyper-parameters that can affect the performance. We fine-tune those parameters on the Pascal Voc 2012 validation set to determine a best-performing combination. Particularly, we set the size of LI zones to be  $3 \times 3$ , the value for the standard deviation  $\sigma$  in Eq. 5 is selected to be 1.0, the LI rate  $e$  in Eq. 3 is set to 1, and all LI intensities  $W_L$  in Eq. 5 are initialised as 0.0 such that the training can start smoothly from any pre-trained checkpoints that do not use LI layers. Moreover, we evaluate different positions of adding LI bottlenecks in MNV2 and ResNet-50 architectures, and a general trend can be observed that adding LI to higher layers can produce better results than to bottom ones. Besides, we select ReLu as the activation  $\phi$  in Eq. 3.

**Implementations** We implement our method in the Tensorflow framework [1]. For the implementation of the baseline Deeplabv3+ [9] model, we directly use the code provided by authors. It takes around one day per GPU (2080TI) to train a model with LI-MNV backbone on Pascal Voc 2012 dataset, and it requires about 2.5/0.6 days to do so on CelebAMask-HQ and ADE20K datasets. For the LI-ResNet-50 backbone, the training will take longer which are approximately 1.5/3.2/4.5 days on Pascal Voc/CelebAMask-HQ/ADE20K using two parallel GPUs.

Backbone	Decoding model	LI Zone Sizes	LI Rates	$W_L$ Init. Range	mIoU (%)
MNV2	ASPP	-	-	-	72.19
MNV2	LI-ASPP	{3, 3, 3}	{1,1,1}	[0.0,0.0]	<b>72.77</b>
		{3, 3, 3}	{1,3,5}	[0.0,0.0]	72.56
		{3, 3, 3}	{5,5,5}	[0.05,0.15]	72.43
		{5, 5, 5}	{1,1,1}	[0.05,0.15]	72.13
		{3, 3, 3}	{1,1,1}	[0.05,0.35]	72.64

Table 1. The performance of different LI-Conv’s parameters for LI-ASPP on Pascal Voc 2012 validation set. The three numbers in “LI Zone Sizes” refers to the sizes of the LI zones (“3” stands for a  $3 \times 3$  LI zone) of the three LI-Conv layers in LI-ASPP, respectively, while “LI Rates” indicates the selections of  $e$  in Eq. 3 for those LI-Conv layers.

Backbone	Decoding model	Positions to add LI	mIoU (%)
MNV2	ASPP	-	72.19
LI-MNV2	ASPP	{1 - 6} <sup>th</sup> RB	72.07
		{16} <sup>th</sup> RB	72.21
		{13, 16} <sup>th</sup> RB	72.43
		{10, 13, 16} <sup>th</sup> RB	<b>72.79</b>

Table 2. The performance when adding LI layers to different positions of MNV2 on Pascal Voc 2012 validation set. “RB” refers to the *Residual Bottleneck* layer in MNV2 [42].

### 4.3. Results

**LI Parameters** In Table 1 we demonstrate the performance of different LI parameters for LI-ASPP (with MNV2 as backbone) on Pascal Voc 2012 validation set. In particular, we investigate the performance of varying settings of LI hyper-parameters such as the size of LI Zones, the LI rates  $e$  and  $W_L$ ’s initialisation range for the three LI-Convs layers in LI-ASPP. As shown in Table 1, most settings can lead to superior performance than the baseline method without any LI-Convs, while using a  $3 \times 3$  LI zone and setting  $e = 1$  can generally yield better performance than other settings like a  $5 \times 5$  LI Zone or  $e = 5$ . LI-ASPP achieves the best performance when all  $W_L$  is initialised from 0.0, potentially due to that the zero initialisation can better encourage a smooth learning of LI intensities, and therefore we opt for this setting for all LI layers.

**Adding LI to MNV2** In addition, we evaluate different options of adding LI-Convs in the *Residual Bottleneck* (RB) layers of the MNV2 architecture [42] on the validation set of Pascal Voc 2012. It can be spotted from Table 2 that adding LI mechanisms to the early RB layers (*e.g.* the earliest six RB layers) cannot promote the accuracy. In contrast, LI-Convs integrated with top layers such as the {10, 13, 16}<sup>th</sup> RB layers can produce higher mIoUs. This observation is somehow in line with the expectations since the higher-level layers are generally encoding more semantic representations, which can better benefit from the improved sensitivity to

Backbone	Decoding model	Positions to add LI	mIoU (%)
ResNet-50	ASPP	-	76.22
LI-ResNet-50	ASPP	conv5_3	76.90
		conv4_6, conv5_3	76.53
		conv3_4	76.21
ResNet-50	LI-ASPP	Three dilated convs	76.94

Table 3. The performance of models with ResNet-50-based backbones on Pascal Voc 2012 validation set. We also explore different positions of adding LI layers to ResNet-50.

semantic contours introduced by LI layers.

**Adding LI to ResNet-50** Table 3 demonstrates the results on Pascal Voc 2012 validation set when adding LI layer to different layers of ResNet-50 architecture with ASPP as the decoding model. We can discover that adding LI to earlier layers of ResNet such as the “conv3\_4” may not improve the performance, however, top layers like “conv4\_6” and “conv5\_3” can better benefit from the integration of LI layers. Such observations are consistent with the trend that is found in the LI-MNV2 experiments of Table 2, which is also in accordance with our intuitions for LI layer’s effects. A slight difference is that the best result is achieved when LI is added to the “conv5\_3” layer other than to both “conv4\_6” and “conv5\_3” layers. Besides, we report in Table 3 the performance of LI-ASPP with ResNet-50 as the backbone, which still shows significant improvement over the baseline.

**Performance Evaluations** In Table 4, we report the evaluation results of different methods with MNV2-based backbones on the three segmentation benchmark datasets. Note that we disable the Deeplabv3+ Decoder [9] in this experiment to ensure a fair and clean comparison. Compared with the baseline method which is MNV2+ASPP, *i.e.* Deeplabv3 [8], LI-MNV2 and LI-ASPP both demonstrate superior performance when used solely, while the best mIoUs on three datasets are all achieved by using them together. Particularly, our method (LI-MNV2+LI-ASPP) gains a relative improvement of 1.32%, 1.30% and 2.30% over the baseline (MNV2+ASPP) on Pascal Voc 2012, CelebAMask-HQ and ADE-20K datasets, respectively, which verifies the effectiveness of LI-Convs. The LI-based model’s parameters and FLOPs, however, are only slightly increased by 0.097% and 0.76% compared with the baseline, which is arguably acceptable considering the accuracy compensations.

The evaluation results of ResNet-50-based models are shown in Table 5, where our methods are additionally compared with Deeplabv3+ Decoder [9], a module that also aims to refine the semantic contours. We can see from the table that when Deeplabv3+ Decoder is disabled, our method (LI-ResNet-50+LI-ASPP) outperforms the baseline (ResNet-50+ASPP) on all three datasets at the cost of slightly increased parameters and FLOPs, which is consistent with

Method	mIoU (%)			Parameters (Kilo)	FLOPs (Mega)
	Pascal Voc 2012	CelebAMask-HQ	ADE-20K		
MNV2 + ASPP	72.19	74.73	29.97	2568.02	6479
MNV2 + LI-ASPP	72.77	75.3	30.47	2568.98	6498
LI-MNV2 + ASPP	72.79	75.46	30.59	2569.94	6517
LI-MNV2 + LI-ASPP	<b>73.14</b>	<b>75.70</b>	<b>30.66</b>	2570.52	6528

Table 4. Performance of different methods with MNV2-based backbones on the Pascal Voc 2012 and ADE20K (validation set) and on the CelebAMask-HQ (test set). The model parameters and FLOPs (for crop size 513 \* 513) are also included.

Method	Deeplabv3+ Decoder [9]	mIoU (%)			Parameters (Kilo)	FLOPs (Giga)
		Pascal Voc 2012	CelebAMask-HQ	ADE-20K		
ResNet-50 + ASPP	-	76.22	76.03	39.14	26656	87.35
LI-ResNet-50 + LI-ASPP	-	77.24	76.62	39.42	26663	87.48
ResNet-50 + ASPP	✓	77.01	77.73	39.87	26819	92.85
LI-ResNet-50 + LI-ASPP	✓	77.54	78.46	40.22	26826	92.98

Table 5. Performance of different methods with ResNet-based backbones on the Pascal Voc 2012 and ADE20K (validation set) and on the CelebAMask-HQ (test set). The performance of disabling/enabling Deeplabv3+ Decoder [9] is reported. The model parameters and FLOPs (for crop size 513 \* 513) are also included.

the MNV2-based results in Table 4. Enabling Deeplabv3+ Decoder introduces mIoU boosts to both our method and the baseline, while our LI models still demonstrate greater improvement over the baseline on all datasets. This indicates that LI-Convs can work closely with Deeplabv3+ Decoder to produce dense predictions with higher-quality, exhibiting the compatibility and the flexibility of integrating LI-Convs into other deep models. Moreover, our LI-based models (LI-ResNet-50+LI-ASPP) with Deeplabv3+ Decoder disabled can achieve similar performance as the baseline (ResNet-50+ASPP) that enables it, while the former model of ours contains 0.58% fewer parameters and operates at approximately 5.78% faster speed than the latter one, respectively, which validates the light-weighted features of the proposed LI-Convs.

**Multi-scale Evaluations** We further compare the performance of different methods when applying the multi-scale evaluation techniques [8, 9]. Particularly, we evaluate results of the LI-based models and the baseline on Pascal Voc 2012 validation set using three different multi-scale settings and with Deeplabv3+ Decoder [9] disabled/enabled. As shown in Table 6, the application of the multi-scale techniques significantly increases the segmentation accuracy of both our and baseline models, while our method consistently outperforms the baseline no matter which multi-scale setting is employed. This is following our expectations, since the proposed LI-Convs can fundamentally enhance the model’s sensitivity to semantic contours, thus will benefit the segmentation results of varying input scales. Additionally, our LI models can work seamlessly with Deeplabv3+ Decoder to achieve the highest mIoUs for all multi-scale settings, which again verifies the generality of LI-Convs.

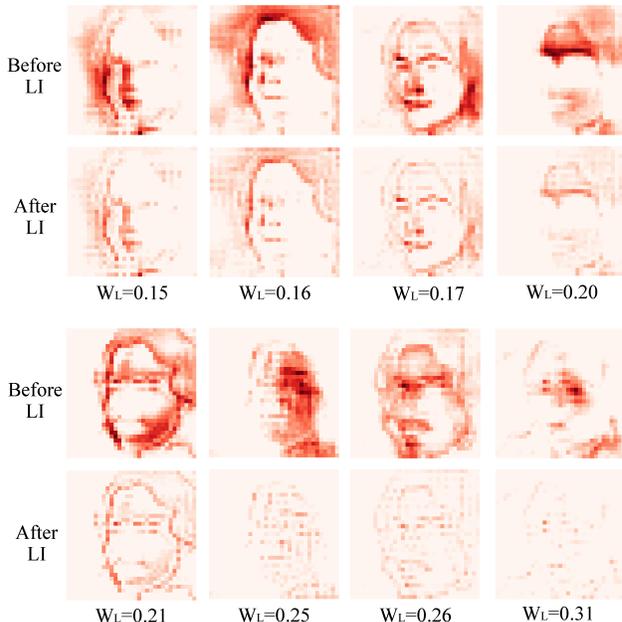


Figure 5. Visualisations of the channel-level features before and after LI layers on CelebAMask-HQ. Although the activation is inhibited globally, the feature patterns after LI layer are generally easier to recognise mainly due to the clarifications on semantic contours.

#### 4.4. Discussion

**How the LI layer works** To intuitively understand the LI mechanisms, we dive into the channel-level features to visualise the patterns before and after LI layers. As demonstrated in Fig. 5, we plot several feature channels before and after the LI layers in LI-ASPP on CelebAMask-HQ dataset.

Method	Deeplabv3+ Decoder [9]	Evaluation Scales			
		[1.0]	[0.5, 1.0, 1.75]	[0.5, 0.75, 1.0, 1.25, 1.75]	[0.5, 0.75, 1.0, 1.25, 1.5, 1.75]
ResNet-50 + ASPP	-	76.22	76.58	77.41	77.60
LI-ResNet-50 + LI-ASPP	-	77.24	77.93	78.37	78.58
ResNet-50 + ASPP	✓	77.01	78.33	78.71	78.72
LI-ResNet-50 + LI-ASPP	✓	77.54	78.66	79.05	79.19

Table 6. Performance (measured by mIoU (%)) of different methods using multi-scale evaluations on Pascal Voc 2012 validation set. “[1.0]” refers to using the single evaluation scale, i.e. no multi-scale is utilised.

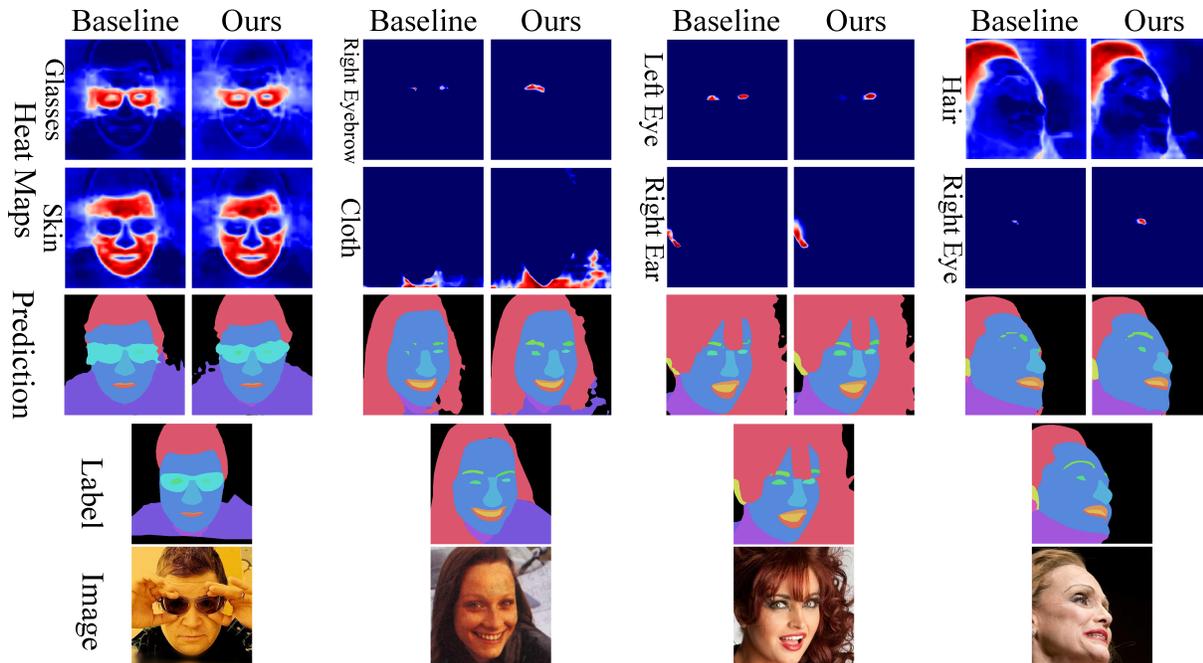


Figure 6. Visualisations of the class-level heat maps and semantic predictions of the baseline (MNv2+ASPP) and our method (LI-MNV2+LI-ASPP) on CelebAMask-HQ. Deeper reds in heat maps represent higher positive responses or more attention from the model, and vice versa for deeper blues. Our method allocates more attention to shape the semantic boundary areas and thus can produce predictions with higher visual qualities.

It can be discovered that although the intensity of activation is suppressed globally after the LI layer, the inhibited feature exhibits more recognisable patterns with clarified and emphasised contours, which can be more desirable in the segmentation domain.

**What interests the model** In Fig. 6, we visualise the class-level heat maps and the segmentation predictions generated by the baseline (MNv2+ASPP) and our method (LI-MNV2+LI-ASPP) on CelebAMask-HQ. We utilise deeper reds to denote higher positive neurons responses (more model attention) in heat maps, and vice versa for deeper blues. Compared with the baseline, the semantically meaningful contouring areas receive more attention from our model, e.g. the “glasses” and “skin” heat maps in Fig. 6. Such kind of contour sensitivity can be reasonably attributed to the proposed LI-Convs. Besides, the segmentation predic-

tions generated by our method have better visual qualities, which also verifies the superiority of the LI-Convs.

## 5. Conclusion

We describe a dilated convolution with lateral inhibitions (LI-Convs) to enhance the model’s sensitivity to semantic contours and to extract features at denser scales. The performance of the proposed LI-ASPP, LI-MNV2 and LI-ResNet architectures is shown to outperform the baseline method on three segmentation benchmark datasets, which verify the effectiveness and generality of the LI-Convs. We also investigate and try to understand the working mechanisms hidden behind. The proposed LI-Convs can be seamlessly integrated into deep models for other tasks, such as lip-reading and object detection, that require explicit awareness of the semantic boundaries.

## References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016. 5
- [2] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(12):2481–2495, 2017. 2
- [3] Chunshui Cao, Yongzhen Huang, Zilei Wang, Liang Wang, Ninglong Xu, and Tieniu Tan. Lateral inhibition-inspired convolutional neural network for visual attention and saliency detection. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018. 3
- [4] Siddhartha Chandra and Iasonas Kokkinos. Fast, exact and multi-scale inference for semantic image segmentation with deep gaussian crfs. In *European Conference on Computer Vision*, pages 402–418. Springer, 2016. 1, 2
- [5] Liang-Chieh Chen, Maxwell Collins, Yukun Zhu, George Papandreou, Barret Zoph, Florian Schroff, Hartwig Adam, and Jon Shlens. Searching for efficient multi-scale architectures for dense image prediction. In *Advances in neural information processing systems*, pages 8699–8710, 2018. 1, 2
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062*, 2014. 1, 2
- [7] 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 transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2018. 1, 2
- [8] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 1, 2, 5, 6, 7
- [9] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018. 1, 2, 4, 5, 6, 7, 8
- [10] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258, 2017. 4, 5
- [11] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 764–773, 2017. 2
- [12] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *International journal of computer vision*, 111(1):98–136, 2015. 2, 4
- [13] Bruno José Torres Fernandes, George DC Cavalcanti, and Tsang Ing Ren. Lateral inhibition pyramidal neural network for image classification. *IEEE transactions on cybernetics*, 43(6):2082–2092, 2013. 2
- [14] Antonio Fernández-Caballero, María T López, Juan Serrano-Cuerda, and José Carlos Castillo. Color video segmentation by lateral inhibition in accumulative computation. *Signal, Image and Video Processing*, 8(6):1179–1188, 2014. 2
- [15] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *2011 International Conference on Computer Vision*, pages 991–998. IEEE, 2011. 4
- [16] Bharath Hariharan, Pablo Arbeláez, Ross Girshick, and Jitendra Malik. Hypercolumns for object segmentation and fine-grained localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 447–456, 2015. 2
- [17] H K Hartline, Henry G Wagner, and Floyd Ratliff. Inhibition in the eye of limulus. *The Journal of general physiology*, 39(5):651–673, 1956. 2
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2, 4, 5
- [19] Matthias Holschneider, Richard Kronland-Martinet, Jean Morlet, and Ph Tchamitchian. A real-time algorithm for signal analysis with the help of the wavelet transform. In *Wavelets*, pages 286–297. Springer, 1990. 1, 2
- [20] Zilong Huang, Xinggang Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. Ccnet: Criss-cross attention for semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 603–612, 2019. 1
- [21] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015. 5
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [23] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017. 2
- [24] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. Maskgan: Towards diverse and interactive facial image manipulation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 4, 5
- [25] Jiqian Li, Yan Wu, Junqiao Zhao, Linting Guan, Chen Ye, and Tao Yang. Pedestrian detection with dilated convolution, region proposal network and boosted decision trees. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pages 4052–4057. IEEE, 2017. 2
- [26] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 5
- [27] Chenxi Liu, Liang-Chieh Chen, Florian Schroff, Hartwig Adam, Wei Hua, Alan L Yuille, and Li Fei-Fei. Auto-deeplab: Hierarchical neural architecture search for semantic image

- segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 82–92, 2019. 1, 2
- [28] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015. 1, 2
- [29] Bingnan Luo, Jie Shen, Shiyang Cheng, Yujiang Wang, and Maja Pantic. Shape constrained network for eye segmentation in the wild. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 1952–1960, 2020. 2
- [30] Pingchuan Ma, Yujiang Wang, Jie Shen, Stavros Petridis, and Maja Pantic. Lip-reading with densely connected temporal convolutional networks. *arXiv preprint arXiv:2009.14233*, 2020. 2
- [31] Zhi-Hong Mao and Steve G Massaquoi. Dynamics of winner-take-all competition in recurrent neural networks with lateral inhibition. *IEEE transactions on neural networks*, 18(1):55–69, 2007. 2
- [32] Brais Martinez, Pingchuan Ma, Stavros Petridis, and Maja Pantic. Lipreading using temporal convolutional networks. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6319–6323. IEEE, 2020. 2
- [33] Shervin Minaee, Yuri Boykov, Fatih Porikli, Antonio Plaza, Nasser Kehtarnavaz, and Demetri Terzopoulos. Image segmentation using deep learning: A survey. *arXiv preprint arXiv:2001.05566*, 2020. 2
- [34] Tuan Nghia Nguyen, Xuan Truong Nguyen, Hyun Kim, and Hyuk-Jae Lee. A lightweight yolov2 object detector using a dilated convolution. In *2019 34th International Technical Conference on Circuits/Systems, Computers and Communications (ITC-CSCC)*, pages 1–2. IEEE, 2019. 2
- [35] Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning deconvolution network for semantic segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1520–1528, 2015. 2
- [36] Chao Peng, Xiangyu Zhang, Gang Yu, Guiming Luo, and Jian Sun. Large kernel matters—improve semantic segmentation by global convolutional network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4353–4361, 2017. 2
- [37] Md Atiqur Rahman and Yang Wang. Optimizing intersection-over-union in deep neural networks for image segmentation. In *International symposium on visual computing*, pages 234–244. Springer, 2016. 2
- [38] Giacomo Rizzolatti and Rosolino Camarda. Inhibition of visual responses of single units in the cat visual area of the lateral suprasylvian gyrus (clare-bishop area) by the introduction of a second visual stimulus. *Brain Research*, 1975. 2
- [39] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 2
- [40] Botond Roska, Erik Nemeth, Laszlo Orzo, and Frank S Werblin. Three levels of lateral inhibition: A space–time study of the retina of the tiger salamander. *Journal of Neuroscience*, 20(5):1941–1951, 2000. 2
- [41] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. 5
- [42] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, 2018. 2, 4, 5, 6
- [43] Falong Shen and Gang Zeng. Gaussian dilated convolution for semantic image segmentation. In *Pacific Rim Conference on Multimedia*, pages 324–334. Springer, 2018. 2
- [44] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3
- [45] C Sun, X Chen, L Huang, and T Shou. Orientation bias of the extraclassical receptive field of the relay cells in the cat’s dorsal lateral geniculate nucleus. *Neuroscience*, 125(2):495–505, 2004. 2
- [46] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8934–8943, 2018. 2
- [47] Georg Von Békésy. *Sensory inhibition*. Princeton University Press, 2017. 2
- [48] Panqu Wang, Pengfei Chen, Ye Yuan, Ding Liu, Zehua Huang, Xiaodi Hou, and Garrison Cottrell. Understanding convolution for semantic segmentation. In *2018 IEEE winter conference on applications of computer vision (WACV)*, pages 1451–1460. IEEE, 2018. 1, 2
- [49] Yujiang Wang, Mingzhi Dong, Jie Shen, Yang Wu, Shiyang Cheng, and Maja Pantic. Dynamic face video segmentation via reinforcement learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6959–6969, 2020. 2
- [50] Yujiang Wang, Bingnan Luo, Jie Shen, and Maja Pantic. Face mask extraction in video sequence. *International Journal of Computer Vision*, 127(6-7):625–641, 2019. 2
- [51] Bo Xu, Cheng Lu, Yandong Guo, and Jacob Wang. Discriminative multi-modality speech recognition. *arXiv preprint arXiv:2005.05592*, 2020. 2
- [52] Shengye Yan, Xinxing Xu, Dong Xu, Stephen Lin, and Xuelong Li. Image classification with densely sampled image windows and generalized adaptive multiple kernel learning. *IEEE transactions on cybernetics*, 45(3):381–390, 2014. 2
- [53] Maoge Yang, Kun Yu, Chi Zhang, Zhiwei Li, and Kuiyuan Yang. Denseaspp for semantic segmentation in street scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3684–3692, 2018. 1, 2
- [54] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015. 1, 2
- [55] Fisher Yu, Vladlen Koltun, and Thomas Funkhouser. Dilated residual networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 472–480, 2017. 1

- [56] Hang Zhang, Kristin Dana, Jianping Shi, Zhongyue Zhang, Xiaoang Wang, Amrith Tyagi, and Amit Agrawal. Context encoding for semantic segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7151–7160, 2018. [1](#)
- [57] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaoang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017. [1](#)
- [58] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 633–641, 2017. [2](#), [4](#), [5](#)
- [59] Yi Zhu and Shawn Newsam. Learning optical flow via dilated networks and occlusion reasoning. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, pages 3333–3337. IEEE, 2018. [2](#)
- [60] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016. [2](#)