

Cars Can't Fly up in the Sky: Improving Urban-Scene Segmentation via Height-driven Attention Networks

Sungha Choi^{1,3} Joanne T. Kim^{2,3} Jaegul Choo⁴

¹A&B Center, LG Electronics, Seoul, South Korea

²Lawrence Livermore National Laboratory, Livermore, CA, USA

³Korea University, Seoul, South Korea ⁴KAIST, Daejeon, South Korea

{shachoi, tengyee}@korea.ac.kr jchoo@kaist.ac.kr

Abstract

This paper exploits the intrinsic features of urban-scene images and proposes a general add-on module, called height-driven attention networks (HANet), for improving semantic segmentation for urban-scene images. It emphasizes informative features or classes selectively according to the vertical position of a pixel. The pixel-wise class distributions are significantly different from each other among horizontally segmented sections in the urban-scene images. Likewise, urban-scene images have their own distinct characteristics, but most semantic segmentation networks do not reflect such unique attributes in the architecture. The proposed network architecture incorporates the capability exploiting the attributes to handle the urban-scene dataset effectively. We validate the consistent performance (mIoU) increase of various semantic segmentation models on two datasets when HANet is adopted. This extensive quantitative analysis demonstrates that adding our module to existing models is easy and cost-effective. Our method achieves a new state-of-the-art performance on the Cityscapes benchmark with a large margin among ResNet-101 based segmentation models. Also, we show that the proposed model is coherent with the facts observed in the urban scene by visualizing and interpreting the attention map. Our code and trained models are publicly available¹.

1. Introduction

Semantic image segmentation, a fundamental task in computer vision, is employed for basic urban-scene understanding in autonomous driving. Fully convolutional networks (FCNs) [28] are seminal work that adopts deep convolutional neural networks (CNNs) in semantic segmentation, by replacing fully connected layers with convolutional ones at the last stage in typical CNN architectures. Other advanced techniques, such as skip-connections in an encoder-decoder architecture [2, 34, 8], an atrous convolu-

tion [6, 46], an atrous spatial pyramid pooling (ASPP) [7], and a pyramid pooling module [51], have further improved the FCN-based architecture in terms of semantic segmentation performance. They have proven to be successful in diverse semantic segmentation benchmarks [14, 15, 31, 12, 1, 25, 32] including urban-scene datasets.

Yet, urban-scene images have their own distinct nature related to perspective geometry [23] and positional patterns [55, 10]. Due to the fact that the urban-scene images are captured by the cameras mounted on the front side of a car, the urban-scene datasets consist only of road-driving pictures. This leads to the possibility of incorporating common structural priors depending on a spatial position, markedly in a vertical position. To verify this characteristic, we present the class distribution of an urban-scene dataset across vertical positions in Fig. 1. Although the pixels of the few classes are dominant in an entire region of an image (Fig. 1(a)), the class distribution has significant dependency on a vertical position. That is, a lower part of an image is mainly composed of road, while the middle part contains various kinds of relatively small objects. In the upper part, buildings, vegetation, and sky are principal objects as shown in Fig. 1(b).

Table 1 shows the probabilities with respect to the dominant top-5 classes: road, building, vegetation, car, and sidewalk. The class distribution is extremely imbalanced that the dominant classes take over 88% of the entire dataset. As mentioned above, the class distribution is completely different if the image is divided into three regions: upper, middle, and lower parts. For instance, the probability of the road class p_{road} is 36.9% on average given an entire image, but this chance drops dramatically to 0.006% for an upper region, while jumps to 87.9% if a lower region is considered.

Also, we analyzed this observation using entropy, a measure of uncertainty. The entropy of the probability distribution X of a pixel over 19 classes in the Cityscapes dataset [12] is computed as

¹<https://github.com/shachoi/HANet>

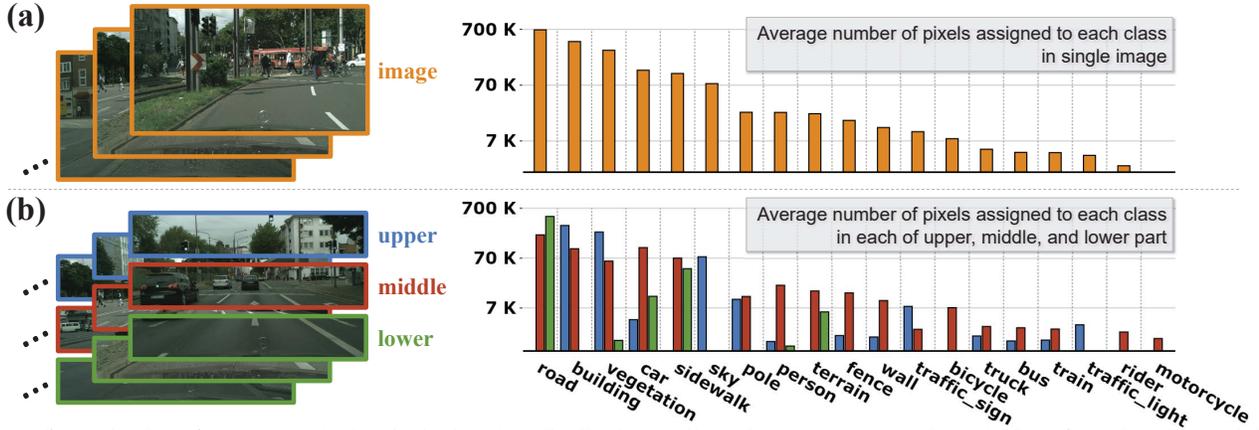


Figure 1: Motivation of our approach, the pixel-wise class distributions. All numbers are average values obtained from the entire training set of the Cityscapes dataset with its pixel-level class labels [12]. Note that there exists a total of 2048K pixels per image, and the y-axis is in log-scale. (a) Each bar represents the average number of pixels assigned to each class contained in a single image. For example, on average, about 685K pixels per image are assigned to the road class. (b) Each part of an image divided into three horizontal sections has a significantly different class distribution from each other. For example, the upper region has just 38 pixels of the road class, while the lower region has 480K pixels of it.

$$\begin{aligned}
 H(X) &= H(p_{\text{road}}, p_{\text{building}}, \dots, p_{\text{motorcycle}}) \\
 &= - \sum_i p_i \log p_i,
 \end{aligned} \tag{1}$$

where p_i denotes the probability that an arbitrary pixel is assigned to the i -th class. the conditional entropy $H(X|\text{image})$ of X given an image, computed by Eq. (1), is 1.84. On the other hand, the average conditional entropy of X given each of the three regions as $H(X|\text{upper})$, $H(X|\text{middle})$, and $H(X|\text{lower})$, is 1.26 as shown in Table 1. As a result, one can see the uncertainty is reduced if we divide an image into several parts horizontally. Based on this analysis, if we can identify the part of an image to which a given arbitrary pixel belongs, it will be helpful for pixel-level classification in semantic segmentation.

Motivated by these observations, we propose a novel height-driven attention networks (HANet) as a general add-on module to semantic segmentation for urban-scene images. Given an input feature map, HANet extracts “height-wise contextual information”, which represents the context of each horizontally divided part, and then predicts which features or classes are more important than the others within each horizontal part from the height-wise contextual information. The models adopting HANet consistently

| Given | Probabilities of top-5 classes | | | | | Entire class entropy | |
|--------|--------------------------------|---------------------|------------------|------------------|--------------------|----------------------|------------|
| | p_{road} | $p_{\text{build.}}$ | p_{veg} | p_{car} | p_{swalk} | | |
| Image | 36.9 | 22.8 | 15.9 | 7.0 | 6.1 | 1.84 | |
| Upper | 0.006 | 47.8 | 35.4 | 0.6 | 0.009 | 1.24 | 1.26 (avg) |
| Middle | 31.4 | 16.6 | 9.4 | 17.4 | 10.7 | 2.04 | |
| Lower | 87.9 | 0.1 | 0.3 | 2.2 | 7.9 | 0.51 | |

Table 1: Comparison of the probability distributions (%) of pixels being assigned to each class when an entire image is separated on upper, middle, and lower regions of the Cityscapes dataset.

tently outperform baseline models. Importantly, our proposed module can be added to any CNN-based backbone networks with negligible cost increase. To verify the effectiveness of HANet, we conduct extensive experiments with various backbone networks such as ShuffleNetV2 [29], MobileNetV2 [37], ResNet-50 [18], and ResNet-101 [18]. We also focus on lightweight backbone networks where a lightweight HANet is more effective. The main contributions of this paper include:

- We propose a novel lightweight add-on module, called HANet, which can be easily added to existing models and improves the performance by scaling the activation of channels according to the vertical position of a pixel. We show the effectiveness and wide applicability of our method through extensive experiments by applying on various backbone networks and two different datasets.
- By adding HANet to the baseline, with negligible computational and memory overhead, we achieve a new state-of-the-art performance on the Cityscapes benchmark with a large margin among ResNet-101 based models.
- We visualize and interpret the attention weights on individual channels to experimentally confirm our intuition and rationale that height position is crucial to improve the segmentation performance on urban scene.

2. Related Work

Model architectures for semantic segmentation. Maintaining the resolution of a feature map while capturing high-level semantic features is essential in achieving high performance of semantic segmentation. Typically, high-level features are extracted by stacking multiple convolutions and spatial pooling layers, but the resolution gets coarser in the process. Several studies [28, 33] address this limitation by leveraging deconvolution for learnable upsampling

from low-resolution features. Skip-connections overcome the limitation by recovering the object boundaries in a decoder layer through leveraging high-resolution features existing earlier in the encoder layer [2, 34, 24, 8]. Another prevalent method is atrous convolution [6, 46], which increases the receptive field size without increasing the number of parameters, and it is widely adopted in recent semantic segmentation networks [7, 8, 43, 49, 53]. Additionally, ASPP [7] and pyramid pooling modules [51] address such challenges caused by diverse scales of objects.

More recently, long-range dependency is captured [54, 21, 16] to improve the performance especially by extending self-attention mechanism [39]. Focusing on boundary information is another approach in semantic segmentation [3, 5, 13, 38, 30]. Recent work imposes separate modules for targeting boundary processing [13, 38] or boundary driven adaptive downsampling [30]. Also, capturing contextual information is widely exploited. ACFNet [48] uses class center to gather features of pixels in each class as categorical context, while CFNet [50] learns the distribution of co-occurrent features and captures co-occurrent context to relook before making predictions. CiSS-Net [52] adopts reinforcement learning to explore the context information in predicted segmentation maps, not having any supervision.

Exploitation of urban-scene image. In the field of semantic segmentation, several studies exploit the characteristics of the urban-scene images. In general, the scale of objects significantly vary in the urban-scene images. FoveaNet [23] localizes a “fovea region”, where the small-scale objects are crowded, and performs scale normalization to address heterogeneous object scales. DenseASPP [43] adopts densely connected ASPP, which connects multiple atrous convolutional layers [20] to address large-scale changes of the objects. Another recent approach [53] exploits the fact that the urban-scene images have continuous video frame sequences and proposes the data augmentation technique based on a video prediction model to create future frames and their labels.

Recent approaches in the field of domain adaptation propose the method to leverage the properties of the urban-scene images. A class-balanced self-training with spatial priors [55] generates pseudo-labels for unlabeled target data by exploiting which classes appear frequently at a particular position in an image for unsupervised domain adaptation. Another approach [10] divides an urban-scene image into several spatial regions and conducts domain adaptation on the pixel-level features from the same spatial region. Also, the correlation between depth information and semantic is exploited to gain additional information from synthetic data for urban-scene domain adaptation [9].

Channel-wise attention. Our proposed method, HANet, has strong connections to a channel-wise attention approach, which exploits the inter-channel relationship of fea-

tures and scales the feature map according to the importance of each channel. Squeeze-and-excitation networks (SENet) [19] capture the global context of the entire image using global average pooling and predict per-channel scaling factors to extract informative features for an image classification task. This mechanism is widely adopted in subsequent studies [40, 49, 44, 22] for image classification and semantic segmentation tasks. Inspired from ParseNet [27], which shows the impact of a global context of an entire image in semantic segmentation, previous work [49, 44, 22] for semantic segmentation exploits the global context of the entire image to generate channel-wise attention. However, the urban-scene datasets [12, 1, 32, 4] consist only of road-driving pictures, which means that the images share similar class statistics. Therefore, the global context should be relatively similar among urban-scene images. As a result, the global context of the entire image cannot present distinct information of each image to help per-pixel classification in urban-scene images. This explains why the previous work related to channel-wise attention for semantic segmentation mainly focuses on the generic scene datasets.

3. Proposed Method

Urban-scene images generally involve common structural priors depending on a spatial position. Each row of an image has significantly different statistics in terms of a category distribution. In this sense, individually capturing the height-wise contextual information, which represents the global context of each row can be used to estimate how channels should be weighted during pixel-level classification for urban-scene segmentation. Therefore, we propose HANet which aims to i) extract the height-wise contextual information and ii) compute height-driven attention weights to represent the importance of features (at intermediate layers) or classes (at last layer) for each row using the context. In this section, we first describe HANet as a general add-on module and then present the semantic segmentation networks incorporating several HANet at different layers specialized for urban-scene segmentation.

3.1. Height-driven Attention Networks (HANet)

HANet generates per-channel scaling factors for each individual row from its height-wise contextual information as its architecture illustrated in Fig. 2.

Let $\mathbf{X}_\ell \in \mathbb{R}^{C_\ell \times H_\ell \times W_\ell}$ and $\mathbf{X}_h \in \mathbb{R}^{C_h \times H_h \times W_h}$ denote the lower- and higher-level feature maps in semantic segmentation networks, where C is the number of channels, H and W are the spatial dimensions of the input tensor, height and width, respectively. Given the lower-level feature map \mathbf{X}_ℓ , F_{HANet} generates a channel-wise attention map $\mathbf{A} \in \mathbb{R}^{C_h \times H_h}$ made up of height-wise per-channel scaling factors and fitted to the channel and height dimensions of the higher-level feature map \mathbf{X}_h . This is done in a series of steps: width-wise pooling (Fig. 2(a)), interpolation for

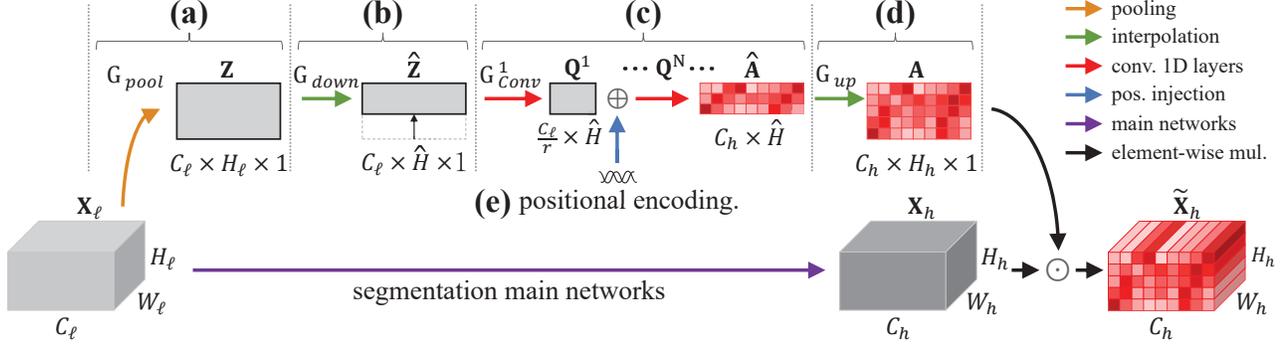


Figure 2: Architecture of our proposed HANet. Each operation op is notated as G_{op} , and feature maps are in bold— X_ℓ : lower-level feature map, Z : width-wise pooled X_ℓ , \hat{Z} : down-sampled Z , Q^n : n -th intermediate feature map of 1D convolution layers, \hat{A} : down-sampled attention map, A : final attention map, X_h : higher-level feature map, \tilde{X}_h : transformed new feature map. Details can be found in Section 3.1.

coarse attention (Fig. 2(b,d)), and computing height-driven attention map (Fig. 2(c)). Moreover, adding positional encoding is included in the process (Fig. 2(e)).

After computing the attention map, the given higher-level feature map X_h can be transformed into a new representation \tilde{X}_h , acquired by an element-wise multiplication of A and X_h . Note that single per-channel scaling vector is derived for each individual row or for each set of multiple consecutive rows, so the vector is copied along with the horizontal direction, which is formulated as

$$\tilde{X}_h = F_{\text{HANet}}(X_\ell) \odot X_h = A \odot X_h. \quad (2)$$

Width-wise pooling (Fig. 2(a)). In order to obtain a channel-wise attention map, we firstly extract height-wise contextual information from each row by aggregating the $C_\ell \times H_\ell \times W_\ell$ input representation X_ℓ into a $C_\ell \times H_\ell \times 1$ matrix Z using a width-wise pooling operation G_{pool} , i.e.,

$$Z = G_{\text{pool}}(X_\ell). \quad (3)$$

There are two typical pooling methods, average pooling and max pooling, to squeeze the spatial dimension. The choice between max pooling and average pooling for the width-wise pooling operation is a hyper-parameter and is empirically set to average pooling. Formally, the h -th row vector of Z is computed as

$$Z_{:,h} = \left[\frac{1}{W} \sum_{i=1}^W \mathbf{X}_{1,h,i}; \dots; \frac{1}{W} \sum_{i=1}^W \mathbf{X}_{C,h,i} \right]. \quad (4)$$

Interpolation for coarse attention (Fig. 2(b,d)). After the pooling operation, the model generates a matrix $Z \in \mathbb{R}^{C_\ell \times H_\ell}$. However, not all the rows of matrix Z may be necessary for computing an effective attention map. As illustrated in Fig. 1(b), class distributions for each part highly differ from each other, even if we divide the entire area into just three parts. Therefore, we interpolate $C_\ell \times H_\ell$ matrix Z into $C_\ell \times \hat{H}$ matrix \hat{Z} by downsampling it (Fig. 2(b)). \hat{H} is a hyper-parameter and is empirically set to 16. Since the attention map, constructed from downsampled representations, is also coarse, the attention map is converted to have

the equivalent height dimension with the given higher-level feature map X_h via upsampling (Fig. 2(d)).

Computation of height-driven attention map (Fig. 2(c)).

A height-driven channel-wise attention map A is obtained by convolutional layers that take width-wise pooled and interpolated feature map \hat{Z} as input. Recent work that utilized a channel-wise attention in classification and semantic segmentation [19, 40, 49] adopts fully connected layers rather than convolutional layers since they generate a channel-wise attention for an entire image. However, we adopt convolutional layers to let the relationship between adjacent rows be considered while estimating the attention map since each row is related to its adjacent rows.

The attention map A indicates which channels are critical at each individual row. There may exist multiple informative features at each row in the intermediate layer; in the last layer, each row can be associated with multiple labels (e.g., road, car, sidewalk, etc.). To allow these multiple features and labels, a sigmoid function is used in computing the attention map, not a softmax function. These operations consisting of N convolutional layers can be written as

$$A = G_{\text{up}} \left(\sigma \left(G_{\text{Conv}}^N \left(\dots \delta \left(G_{\text{Conv}}^1(\hat{Z}) \right) \right) \right) \right), \quad (5)$$

where σ is a sigmoid function, δ is a ReLU activation, and G_{Conv}^i denotes i -th one-dimensional convolutional layer. We empirically adopt three convolutional layers: the first one $G_{\text{Conv}}^1(\hat{Z}) = Q^1 \in \mathbb{R}^{\frac{C_\ell}{r} \times \hat{H}}$ for channel reduction, the second one $G_{\text{Conv}}^2(\delta(Q^1)) = Q^2 \in \mathbb{R}^{2 \cdot \frac{C_\ell}{r} \times \hat{H}}$, and the last one $G_{\text{Conv}}^3(\delta(Q^2)) = \hat{A} \in \mathbb{R}^{C_h \times \hat{H}}$ for generating an attention map. The reduction ratio r reduces the parameter overhead of HANet as well as gives a potential regularization effect. An analysis on the effect of various reduction ratio as a hyper-parameter will be presented in Section 4.2.2.

Incorporating positional encoding (Fig. 2(e)). When humans recognize a driving scene, they have prior knowledge on the vertical position of particular objects (e.g., road and sky appear in the lower and upper part, respectively). Inspired by this observation, we add the sinusoidal positional

encodings [39] to the intermediate feature map Q^i at the i -th layer in the HANet. A hyper-parameter i is analyzed in the supplementary material. For injecting positional encodings, we follow the strategy proposed in Transformer [39]. The dimension of the positional encodings is same as the channel dimension C of the intermediate feature map Q^i . The positional encodings are defined as

$$PE_{(p,2i)} = \sin(p/100^{2i/C})$$

$$PE_{(p,2i+1)} = \cos(p/100^{2i/C}),$$

where p denotes the vertical position index in the entire image ranging from zero to $\hat{H} - 1$ of coarse attention, and i is the dimension. The number of the vertical position is set to \hat{H} as the number of rows in coarse attention. The new representation \tilde{Q} incorporating positional encodings is formulated as

$$\tilde{Q} = Q \oplus PE, \quad (6)$$

where \oplus is an element-wise sum.

Height positions are randomly jittered by up to two positions to generalize over different camera location from various datasets to prevent an inordinately tight position-object coupling. Additionally, we experiment with using learnable positional embeddings [17] to find the best way to incorporate positional information in the supplementary material.

Meanwhile, CoordConv [26] proposed to embed height and width coordinates in the intermediate features for various vision tasks: extra channels containing hard-coded coordinates (e.g., height, width, and optional r) are concatenated channel-wise to the input representation, and then a standard convolutional layer is applied. Unlike this model, HANet exploits positional information of height to obtain attention values, $C \times H \times 1$, which is used for gating the output representations of main networks. Therefore, HANet differs significantly from CoordConv in terms of how to exploit the positional information. We experimentally compare ours with CoordConv in Section 4.2.2.

3.2. Segmentation Networks based on HANet

We adopt DeepLabv3+ [6] as a baseline for semantic segmentation. DeepLabv3+ has an encoder-decoder architecture with ASPP that employs various dilation rates. We add HANet to the segmentation networks at five different layers (Fig. 3) after the point where high-level representation is encoded from backbone networks. This is because the higher-level feature has a stronger correlation with the vertical position. We conduct an ablation study to see the performance gain from adding HANet at different layers.

3.3. Comparison with Other Attention Strategies

Self-attention-based approaches like DANet [16], obtain attention values, $(H \times W) \times (H \times W)$, $C \times C$, from the semantic interdependencies in spatial and channel dimensions, respectively. However, HANet does not consider

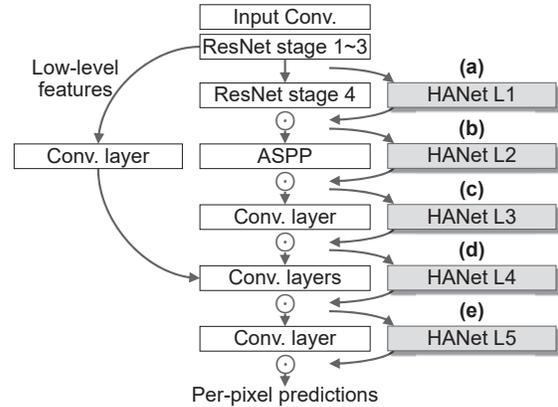


Figure 3: Semantic segmentation networks incorporating HANet in five different layers.

the interdependencies among the dimensions. HANet includes one-dimensional convolutional layers being a separate branch as a modular design (Fig. 2(a)-(d)), which specifically consider the structural property in urban-scene data. In this manner, HANet derives attention values, $C \times H \times 1$, to gate activation at the horizontal section of a feature map output in the main networks, considering vertical position. HANet is significantly more lightweight than self-attention-based approaches that considers the relationship between every pair in each dimension. Meanwhile, channel-wise attention approaches such as SENet [19] generates attention values, $C \times 1 \times 1$, only at an image level. This is not ideal for urban-scene segmentation as most of the images shares similar circumstances and context.

4. Experiments

In this section, we first describe the implementation details of HANet. Then, we experimentally demonstrate the effectiveness and wide applicability of our proposed methods by extensive quantitative analysis including ablation studies. We evaluate HANet on two different urban-scene datasets including Cityscapes [12] and BDD100K [45]. Furthermore, we visualize and analyze the attention map generated from HANet. For all the quantitative experiments, we measure the segmentation performance in terms of mean Intersection over Union (mIoU) metric.

4.1. Experimental Setup

Base segmentation architecture Our network architecture for semantic segmentation is based on DeepLabv3+ [8]. We adopt various backbone networks including ShuffleNetV2 [29], MobileNetV2 [37], and ResNet [18] to verify wide applicability of HANet. Note that HANet can be easily inserted on top of various backbone networks. The adopted backbone networks are pretrained on ImageNet [36] for all the experiments, unless otherwise noted.

Stronger baseline To strictly verify the effectiveness of HANet, we improved the performance of the DeepLabv3+ baseline adopting ResNet-101, by integrating SyncBatch-

| Backbone | OS | Models | Params | GFLOPs | mIoU(%) |
|-------------------------|----|----------|--------|---------|--------------|
| ShuffleNet V2 (1×) [29] | 32 | Baseline | 12.6M | 64.34 | 70.27 |
| | | +HANet | 14.9M | 64.39 | 71.30 |
| | 16 | Baseline | 12.6M | 117.09 | 70.85 |
| | | +HANet | 13.7M | 117.14 | 71.52 |
| MobileNet V2 [29] | 16 | Baseline | 14.8M | 142.74 | 73.93 |
| | | +HANet | 16.1M | 142.80 | 74.96 |
| | 8 | Baseline | 14.8M | 428.70 | 73.40 |
| | | +HANet | 15.4M | 428.82 | 74.70 |
| ResNet-50 [18] | 16 | Baseline | 45.1M | 553.74 | 76.84 |
| | | +HANet | 47.6M | 553.85 | 77.78 |
| | 8 | Baseline | 45.1M | 1460.56 | 77.76 |
| | | +HANet | 46.3M | 1460.76 | 78.71 |
| ResNet-101 [18] | 16 | Baseline | 64.2M | 765.53 | 77.80 |
| | | +HANet | 65.4M | 765.63 | 79.31 |
| | 8 | Baseline | 64.2M | 2137.82 | 79.25 |
| | | +HANet | 65.4M | 2138.02 | 80.29 |

Table 2: Comparison of mIoU, the number of model parameters and FLOPs between the baseline and HANet on Cityscapes validation set according to various backbone networks and output stride. Adding HANet to the baseline consistently increase the mIoU with minimal cost increase.

Norm (batch statistics synchronized across multiple GPUs) publicly included in PyTorch v1.1 and by replacing a single 7×7 convolution by three 3×3 convolutions in the first layer of ResNet-101. we also adopted an auxiliary cross-entropy loss in the intermediate feature map and class uniform sampling [35, 53] to handle class imbalance problems. As a result, our baseline achieves the mIoU of 79.25% on the Cityscapes validation set, which surpasses the other baseline models based on DeepLabv3+ architecture with ResNet-101 of previous work.

Training protocol. We employ SGD optimizer with initial learning rate of $1e-2$ and momentum of 0.9. The weight decays are $5e-4$ and $1e-4$ for main networks and HANet, respectively. The learning rate scheduling follows the polynomial learning rate policy [27]. The initial learning rate is multiplied by $(1 - \frac{iteration}{max.iteration})^{power}$, where power is 0.9. To avoid overfitting, typical data augmentations in semantic image segmentation models are used, including random horizontally flipping, random scaling in the range of [0.5,2], gaussian blur, color jittering, and random cropping.

4.2. Cityscapes

The Cityscapes dataset [12] is a large-scale urban-scene dataset, holding high-quality pixel-level annotations of 5K images and 20K coarsely annotated images. Finely annotated images consist of 2,975 train images, 500 validation images, and 1,525 test images. The annotations of test images are withheld for benchmarks. The resolution of each image is 2048×1024 , and 19 semantic labels are defined. In all the experiments on Cityscapes validation set, we train our models using finely annotated training set for 40K iterations with a total batch size of 8 and a crop size of 768×768 .

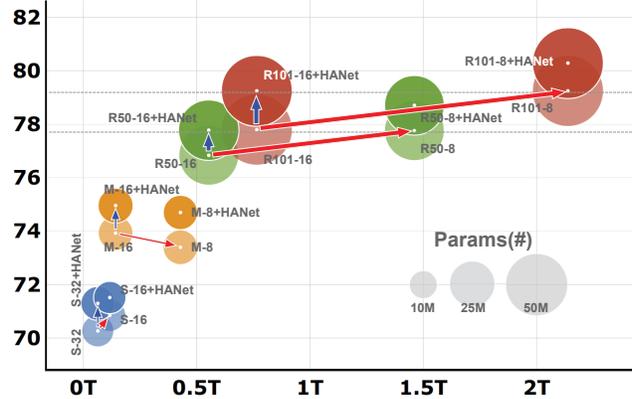


Figure 4: Comparison of the performance and complexity among the baseline and HANet on the various backbone networks. x-axis denotes teraFLOPs and y-axis denotes mIoU. The circle size denotes the number of model parameters. The texts in the colored circle indicate backbone networks, output stride, and whether HANet is adopted to the baseline. S, M, R50, and R101 denote ShuffleNetV2, MobileNetV2, ResNet-50, and -101, respectively. (e.g., S-16: Baseline, ShuffleNetV2, and output stride 16)

4.2.1 Effectiveness of the HANet components.

Table 2 shows the effect of adopting HANet through performance increase (mIoU) according to the number of parameters and FLOPs which indicate model size and complexity, respectively. To demonstrate the wide applicability of HANet, various backbones are examined including ShuffleNetV2 [29], MobileNetV2 [37], ResNet-50, and ResNet-101 [18]. Models with HANet consistently outperform baseline models with significant increases on MobileNetV2 and ResNet-101. Moreover, the model parameters and complexity results indicate that the cost of adding HANet is practically negligible. From Fig. 4, we can see clearly that adding HANet (blue arrow) is worth more than it costs in FLOPs, compared to improving the model through changing output stride (red arrow). Therefore, HANet has a great advantage of not only an effective way of improving semantic segmentation accuracy, but also lightweight algorithm design for practical usage.

4.2.2 Ablation studies

For ablation studies, we use ResNet-101 backbone with output stride of 16 and evaluate on the Cityscapes validation set. Table 3 shows the performance gain when HANet is added to the multiple layers and incorporates the positional encodings. Additionally, we conduct experiments by changing channel reduction ratio r and pooling method. When we add HANet including positional encodings at multiple layers from L1 to L4 (Fig. 3) and the reduction ratio is set to 32, the mIoU significantly increases from 77.80 to 79.31. To compare with CoordConv [26], we conduct experiments by replacing standard convolutional layers after the backbone with CoordConv. Table 3 shows that HANet outperforms the baseline with CoordConv.

| Layers | | | | | Positional encoding | Ratio r | Pooling method | mIoU |
|--|---|---|---|---|---------------------|-----------|----------------|--------------|
| 1 | 2 | 3 | 4 | 5 | | | | |
| ✓ | | | | | ✓ | 32 | average | 78.52 |
| ✓ | ✓ | | | | ✓ | 32 | average | 78.85 |
| ✓ | ✓ | ✓ | | | ✓ | 32 | average | 78.72 |
| ✓ | ✓ | ✓ | ✓ | | ✓ | 32 | average | 79.31 |
| ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | 32 | average | 78.79 |
| ✓ | ✓ | ✓ | ✓ | | ✓ | 16 | average | 79.15 |
| ✓ | ✓ | ✓ | ✓ | | ✓ | 64 | average | 79.08 |
| ✓ | ✓ | ✓ | ✓ | | | 32 | average | 78.25 |
| ✓ | ✓ | ✓ | ✓ | | ✓ | 32 | max | 78.87 |
| Baseline | | | | | | | | 77.80 |
| Baseline + CoordConv [26] (Height + Width) | | | | | | | | 78.82 |

Table 3: Ablation studies and hyper-parameter impacts with regard to the HANet injected layers, using positional encodings or not, and channel reduction ratio. ResNet-101, output stride 16 on Cityscapes validation set.

4.2.3 Inference techniques.

For further performance improvements, we adopt frequently used techniques such as left-right flipped, multi-scale (with $scales=\{0.5, 1.0, 2.0\}$) and sliding inference. In such manner, our best model achieves 82.05% mIoU on the Cityscapes validation set as in Table 4.

| Inference techniques | Baseline | +HANet |
|-----------------------------------|----------|--------------|
| None | 79.25 | 80.29 |
| Multiscale, Sliding, and Flipping | 81.14 | 82.05 |

Table 4: mIoU(%) comparison with respect to inference techniques. ResNet-101, output stride 8 on Cityscapes validation set.

4.2.4 Efficacy at segmented regions

As mentioned in Section 1, the average entropy decreases as we divide the image into multiple horizontal subsections. This implies the performance improvement in the entire region of an image. Besides, the entropy of the upper and lower regions have low entropy compared to the middle region. In this respect, we expect the performance increase arising from HANet would be larger in the upper and lower regions than that in the middle or entire region. Indeed, the performance significantly rises on the upper and lower regions as in Table 5.

| Model (Year) | Backbone | mIoU | road | swalk | build. | wall | fence | pole | tligh. | tsign | veg | terr. | sky | pers. | rider | car | truck | bus | train | meyc | beyc |
|----------------------------------|-------------------|-------------|------|-------|--------|------|-------|------|--------|-------|------|-------|------|-------|-------|------|-------|------|-------|------|------|
| BFP [13]([†] 19) | ResNet-101 | 81.4 | 98.7 | 87.0 | 93.5 | 59.8 | 63.4 | 68.9 | 76.8 | 80.9 | 93.7 | 72.8 | 95.5 | 87.0 | 72.1 | 96.0 | 77.6 | 89.0 | 86.9 | 69.2 | 77.6 |
| DANet [16]([†] 19) | ResNet-101 | 81.5 | 98.6 | 86.1 | 93.5 | 56.1 | 63.3 | 69.7 | 77.3 | 81.3 | 93.9 | 72.9 | 95.7 | 87.3 | 72.9 | 96.2 | 76.8 | 89.4 | 86.5 | 72.2 | 78.2 |
| CCNet [16]([†] 19) | ResNet-101 | 81.5 | 98.6 | 86.1 | 93.5 | 56.1 | 63.3 | 69.7 | 77.3 | 81.3 | 93.9 | 72.9 | 95.7 | 87.3 | 72.9 | 96.2 | 76.8 | 89.4 | 86.5 | 72.2 | 78.2 |
| ACFNet [48]([†] 19) | ResNet-101 | 81.8 | 98.7 | 87.1 | 93.9 | 60.2 | 63.9 | 71.1 | 78.6 | 81.5 | 94.0 | 72.9 | 95.9 | 88.1 | 74.1 | 96.5 | 76.6 | 89.3 | 81.5 | 72.1 | 79.2 |
| Ours | ResNet-101 | 82.1 | 98.8 | 88.0 | 93.9 | 60.5 | 63.3 | 71.3 | 78.1 | 81.3 | 94.0 | 72.9 | 96.1 | 87.9 | 74.5 | 96.5 | 77.0 | 88.0 | 85.9 | 72.7 | 79.0 |
| DeepLabv3+ [8]([†] 18) | Xception [11] | 82.1 | 98.7 | 87.0 | 93.9 | 59.5 | 63.7 | 71.4 | 78.2 | 82.2 | 94.0 | 73.0 | 95.8 | 88.0 | 73.3 | 96.4 | 78.0 | 90.9 | 83.9 | 73.8 | 78.9 |
| GSCNN [38]([‡] 19) | WideResNet38 [47] | 82.8 | 98.7 | 87.4 | 94.2 | 61.9 | 64.6 | 72.9 | 79.6 | 82.5 | 94.3 | 74.3 | 96.2 | 88.3 | 74.2 | 96.0 | 77.2 | 90.1 | 87.7 | 72.6 | 79.4 |
| Ours ^{†‡} | ResNext-101 [42] | 83.2 | 98.8 | 88.0 | 94.2 | 66.6 | 64.8 | 72.0 | 78.2 | 81.4 | 94.2 | 74.5 | 96.1 | 88.1 | 75.6 | 96.5 | 80.3 | 93.2 | 86.6 | 72.5 | 78.7 |

Table 6: Comparison of mIoU and per-class IoU with other state-of-the-art models on Cityscapes test set. [†] denotes training including Cityscapes coarsely annotated images. [‡] denotes training with Mapillary pretrained model.

| Model | Upper | Mid-upper | Mid-lower | Lower | Entire |
|-------------|--------------|-----------|-----------|--------------|--------|
| Baseline | 78.69 | 76.35 | 83.16 | 70.59 | 81.14 |
| +HANet | 80.29 | 77.09 | 84.09 | 73.04 | 82.05 |
| Increase(%) | +1.60 | +0.74 | +0.93 | +2.45 | +0.91 |

Table 5: mIoU(%) comparison to baseline on each part of image divided into four horizontal sections. ResNet-101, output stride 8 on Cityscapes validation set.

4.2.5 Comparison to other state-of-the-art models

To compare with other state-of-the-arts, we train our models using finely annotated training and validation set for 90K iterations. In case of adopting ResNext-101 [42] backbone, coarsely annotated images are additionally used and the model is pretrained on Mapillary [32]. The crop size and batch size are changed into 864×864 and 12, respectively. The inference techniques from Section 4.2.3 are used, but we do not adopt any other techniques such as online bootstrapping for hard training pixels [41]. We compare our best models based on ResNet-101 and ResNext-101 with other recent models on the Cityscapes test set (Table 6). Our models achieve a new state-of-the-art performance.

4.3. BDD100K

Berkeley Deep Drive dataset (BDD100K) [45] is a large-scale diverse driving video database. It includes a semantic image segmentation dataset, consisting of 7,000 training and 1,000 validation images with a resolution of 1280×720 . It is a challenging dataset including images of various driving circumstance such as day, night, and diverse weather conditions. Table 7 shows the superior results of HANet in BDD100K. The training strategy of BDD100K is similar to Cityscapes, but we change the crop size into 608×608 and train for 60K iterations with a total batch size of 16. These adjustments simply comes from smaller image size and larger dataset size compared to Cityscapes.

4.4. Qualitative Analysis

Attention map visualization. We visualize the attention weights to analyze the behavior of the proposed HANet. The attention visualization highlights those channels emphasized at a different vertical position. Through the visualization, we can find out interesting clues to validate our

| Backbone | OS | Baseline | +HANet |
|-------------|----|----------|---------------|
| MobileNetV2 | 16 | 58.91% | 60.05% |
| ResNet-101 | 16 | 63.75% | 64.56% |
| ResNet-101 | 8 | 64.84% | 65.60% |

Table 7: Comparison of mIoU between the baseline and HANet on BDD100K validation set. The output stride is set to 16 and the crop size is 608×608 .

observations and methods.

Fig. 5 clearly shows that HANet assigns a different amount of attention to a different vertical position, indicating that the model properly learns structural priors with respect to the height in urban-scene data.

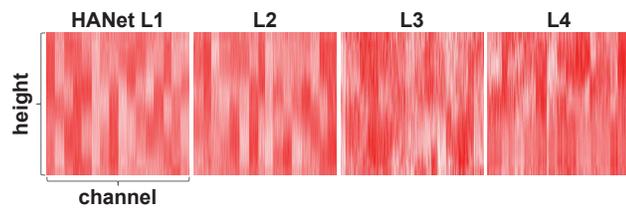


Figure 5: Visualization of attention map from HANet at different layers. x-axis denotes the channels, and y-axis denotes the height of the feature map, showing which channels are weighted at different vertical position. The channels focused by each height are clearly different. To better visualize, the channels are clustered.

Fig. 6 visualizes the attention map from HANet L2, which computes the attention weights for ASPP layer. In ASPP layer, the channels are obtained from convolution filters that have multi-scale receptive fields and grouped for each scale. Therefore, we can interpret the HANet L2 attention map by each group, with the sequence of channels remaining unchanged. Colored boxes give us insights of our method and the urban-scene images. The green boxes in Fig. 6 shows the low-focused channels of the middle region of images, while blue box indicates the channels which are relatively more focused. That is, the channels obtained from the small receptive field are weighted in the middle region. Since the middle region is where small objects are crowded as pointed in the left figure in Fig. 6, small receptive fields

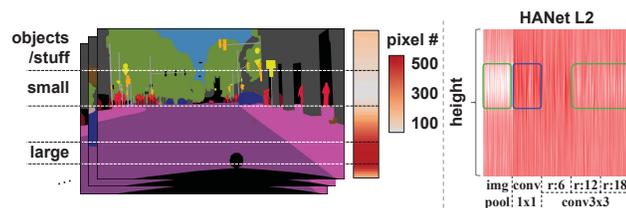


Figure 6: (left) The color in the heatmap denotes an average number of pixels that make up an object. Relatively small objects are crowded in the middle regions, while large objects exist in the lower region. (right) Visualization of attention map from HANet at the second layer. Unlike Fig. 5, the sequence of channels remains unchanged.

are effective for this region and vice versa. In this manner, we verify that HANet properly learns and captures the intrinsic features of urban scene.

Fig. 7 illustrates that the distribution of the attention map (right figure) from HANet at the last layer, which is following the actual height-wise class distribution (left figure) obtained from the Cityscapes training images. Each class gives a different weight according to the vertical position, meaning that the model actually uses vertical positions in the image. This information corresponds to the observation we introduced through the class distribution analysis in Fig. 1 and Table 1. For instance, road (class 0) appears only in the middle and the lower regions, while sky (class 10) is mainly emphasized in the upper rows.

Visualization of semantic segmentation results are shown in the supplementary material.

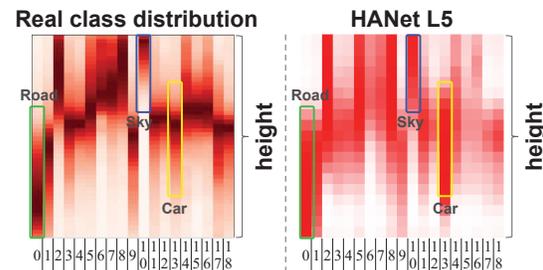


Figure 7: Height-wise class distributions and attention map visualization (L5). The number ranging from 0 to 18 indicates a different class. The darker it is colored, the higher probability (more pixels) assigned to a particular class. The attention visualization follows the patterns in the real class distribution.

5. Conclusions

In this paper, we proposed HANet for urban-scene segmentation as an effective and wide applicable add-on module. Our method exploits the spatial priors existing in urban-scene images to construct a cost-effective architecture. We demonstrated the performance increase by adding our method to the baseline model with negligible cost increase. Moreover, we visualized and analyzed the attention maps to show the validity of our initial hypothesis that exploiting vertical positional information helps for urban-scene semantic segmentation.

Acknowledgments This work was partially supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2019-0-00075, Artificial Intelligence Graduate School Program (KAIST)), the National Research Foundation of Korea (NRF) grant funded by the Korean government (MSIP) (No. NRF-2018M3E3A1057305), and the National Supercomputing Center with supercomputing resources including technical support (KSC-2019-INO-0011).

References

- [1] Hassan Alhaija, Siva Mustikovela, Lars Mescheder, Andreas Geiger, and Carsten Rother. Augmented reality meets computer vision: Efficient data generation for urban driving scenes. *International Journal of Computer Vision (IJCV)*, 2018. 1, 3
- [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 (TPAMI)*, 39(12):2481–2495, 2017. 1, 3
- [3] Gedas Bertasius, Jianbo Shi, and Lorenzo Torresani. High-for-low and low-for-high: Efficient boundary detection from deep object features and its applications to high-level vision. In *Proc. of the International Conference on Computer Vision (ICCV)*, pages 504–512, 2015. 3
- [4] Gabriel J. Brostow, Jamie Shotton, Julien Fauqueur, and Roberto Cipolla. Segmentation and recognition using structure from motion point clouds. In *ECCV*, pages 44–57, 2008. 3
- [5] Liang-Chieh Chen, Jonathan T Barron, George Papandreou, Kevin Murphy, and Alan L Yuille. Semantic image segmentation with task-specific edge detection using cnns and a discriminatively trained domain transform. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4545–4554, 2016. 3
- [6] 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 (TPAMI)*, 40(4):834–848, 2018. 1, 3, 5
- [7] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 1, 3
- [8] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proc. of the European Conference on Computer Vision (ECCV)*, pages 801–818, 2018. 1, 3, 5, 7
- [9] Yuhua Chen, Wen Li, Xiaoran Chen, and Luc Van Gool. Learning semantic segmentation from synthetic data: A geometrically guided input-output adaptation approach. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1841–1850, 2019. 3
- [10] Yuhua Chen, Wen Li, and Luc Van Gool. Road: Reality oriented adaptation for semantic segmentation of urban scenes. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7892–7901, 2018. 1, 3
- [11] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1251–1258, 2017. 7
- [12] 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 *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3213–3223, 2016. 1, 2, 3, 5, 6
- [13] Henghui Ding, Xudong Jiang, Ai Qun Liu, Nadia Magnenat Thalmann, and Gang Wang. Boundary-aware feature propagation for scene segmentation. In *Proc. of the International Conference on Computer Vision (ICCV)*, pages 6819–6829, 2019. 3, 7
- [14] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision (IJCV)*, 88(2):303–338, 2010. 1
- [15] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>. 1
- [16] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3146–3154, 2019. 3, 5, 7
- [17] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. Convolutional sequence to sequence learning. In *Proc. the International Conference on Machine Learning (ICML)*, pages 1243–1252. JMLR. org, 2017. 5
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 2, 5, 6
- [19] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7132–7141, 2018. 3, 4, 5
- [20] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4700–4708, 2017. 3
- [21] Zilong Huang, Xinggang Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. Ccnet: Criss-cross attention for semantic segmentation. In *Proc. of the International Conference on Computer Vision (ICCV)*, pages 603–612, 2019. 3
- [22] Hanchao Li, Pengfei Xiong, Jie An, and Lingxue Wang. Pyramid attention network for semantic segmentation. In *British Machine Vision Conference (BMVC)*, 2018. 3
- [23] Xin Li, Zequn Jie, Wei Wang, Changsong Liu, Jimei Yang, Xiaohui Shen, Zhe L. Lin, Qiang Chen, Shuicheng Yan, and Jiashi Feng. Foveanet: Perspective-aware urban scene parsing. In *Proc. of the International Conference on Computer Vision (ICCV)*, pages 784–792, 2017. 1, 3
- [24] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1925–1934, 2017. 3
- [25] 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 *Proc. of the European Conference on Computer Vision (ECCV)*, pages 740–755. Springer, 2014. 1
- [26] Rosanne Liu, Joel Lehman, Piero Molino, Felipe Petroski Such, Eric Frank, Alex Sergeev, and Jason Yosinski. An intriguing failing of convolutional neural networks and the coordconv solution. In *Proc. the Advances in Neural Information Processing Systems (NeurIPS)*, pages 9605–9616, 2018. 5, 6, 7
- [27] Wei Liu, Andrew Rabinovich, and Alexander C Berg. Parsenet: Looking wider to see better. In *CoRR*. Citeseer, 2015. 3, 6
- [28] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3431–3440, 2015. 1, 2
- [29] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Proc. of the European Conference on Computer Vision (ECCV)*, pages 116–131, 2018. 2, 5, 6
- [30] Dmitrii Marin, Zijian He, Peter Vajda, Priyam Chatterjee, Sam Tsai, Fei Yang, and Yuri Boykov. Efficient segmentation: Learning downsampling near semantic boundaries. In *Proc. of the International Conference on Computer Vision (ICCV)*, pages 2131–2141, 2019. 3
- [31] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. The role of context for object detection and semantic segmentation in the wild. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 891–898, 2014. 1
- [32] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Bulò, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *Proc. of the International Conference on Computer Vision (ICCV)*, pages 4990–4999, 2017. 1, 3, 7
- [33] Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning deconvolution network for semantic segmentation. In *Proc. of the International Conference on Computer Vision (ICCV)*, pages 1520–1528, 2015. 2
- [34] 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. 1, 3
- [35] Samuel Rota Bulò, Lorenzo Porzi, and Peter Kotschieder. In-place activated batchnorm for memory-optimized training of dnns. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5639–5647, 2018. 6
- [36] 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 (IJCV)*, 115(3):211–252, 2015. 5
- [37] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proc. of the European Conference on Computer Vision (ECCV)*, Jun 2018. 2, 5, 6
- [38] Towaki Takikawa, David Acuna, Varun Jampani, and Sanja Fidler. Gated-scnn: Gated shape cnns for semantic segmentation. In *Proc. of the International Conference on Computer Vision (ICCV)*, pages 5229–5238, 2019. 3, 7
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proc. the Advances in Neural Information Processing Systems (NeurIPS)*, pages 5998–6008, 2017. 3, 5
- [40] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *Proc. of the European Conference on Computer Vision (ECCV)*, pages 3–19, 2018. 3, 4
- [41] Zifeng Wu, Chunhua Shen, and Anton van den Hengel. High-performance semantic segmentation using very deep fully convolutional networks, 2016. 7
- [42] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1492–1500, 2017. 7
- [43] Maoke Yang, Kun Yu, Chi Zhang, Zhiwei Li, and Kuiyuan Yang. Denseaspp for semantic segmentation in street scenes. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3684–3692, 2018. 3
- [44] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Learning a discriminative feature network for semantic segmentation. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1857–1866, 2018. 3
- [45] Fisher Yu. Bdd100k: A diverse driving video database with scalable annotation tooling. *arXiv*, 2018. 5, 7
- [46] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015. 1, 3
- [47] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016. 7
- [48] Fan Zhang, Yanqin Chen, Zhihang Li, Zhibin Hong, Jingtuo Liu, Feifei Ma, Junyu Han, and Errui Ding. Acfnnet: Attentional class feature network for semantic segmentation. In *Proc. of the International Conference on Computer Vision (ICCV)*, 2019. 3, 7
- [49] Hang Zhang, Kristin Dana, Jianping Shi, Zhongyue Zhang, Xiaoang Wang, Amrbrish Tyagi, and Amit Agrawal. Context encoding for semantic segmentation. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3, 4
- [50] Hang Zhang, Han Zhang, Chenguang Wang, and Junyuan Xie. Co-occurrent features in semantic segmentation. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 548–557, 2019. 3
- [51] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaoang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2881–2890, 2017. 1, 3
- [52] Yizhou Zhou, Xiaoyan Sun, Zheng-Jun Zha, and Wenjun Zeng. Context-reinforced semantic segmentation. In *Proc.*

of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 4046–4055, 2019. 3

- [53] Yi Zhu, Karan Sapra, Fitsum A Reda, Kevin J Shih, Shawn Newsam, Andrew Tao, and Bryan Catanzaro. Improving semantic segmentation via video propagation and label relaxation. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8856–8865, 2019. 3, 6
- [54] Zhen Zhu, Mengde Xu, Song Bai, Tengeng Huang, and Xiang Bai. Asymmetric non-local neural networks for semantic segmentation. In *Proc. of the International Conference on Computer Vision (ICCV)*, pages 593–602, 2019. 3
- [55] Yang Zou, Zhiding Yu, BVK Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proc. of the European Conference on Computer Vision (ECCV)*, pages 289–305, 2018. 1, 3