



ESPNet: Efficient Spatial Pyramid of Dilated Convolutions for Semantic Segmentation

Sachin Mehta¹ , Mohammad Rastegari², Anat Caspi¹, Linda Shapiro¹, and Hannaneh Hajishirzi¹

¹ University of Washington, Seattle, WA, USA

{sacmehta,caspian,shapiro,hannaneh}@cs.washington.edu

² Allen Institute for AI and XNOR.AI, Seattle, WA, USA
mohammadr@allenai.org

Abstract. We introduce a fast and efficient convolutional neural network, ESPNet, for semantic segmentation of high resolution images under resource constraints. ESPNet is based on a new convolutional module, efficient spatial pyramid (ESP), which is efficient in terms of computation, memory, and power. ESPNet is 22 times faster (on a standard GPU) and 180 times smaller than the state-of-the-art semantic segmentation network PSPNet, while its category-wise accuracy is only 8% less. We evaluated ESPNet on a variety of semantic segmentation datasets including Cityscapes, PASCAL VOC, and a breast biopsy whole slide image dataset. Under the same constraints on memory and computation, ESPNet outperforms all the current efficient CNN networks such as MobileNet, ShuffleNet, and ENet on both standard metrics and our newly introduced performance metrics that measure efficiency on edge devices. Our network can process high resolution images at a rate of 112 and 9 frames per second on a standard GPU and edge device, respectively. Our code is open-source and available at <https://sacmehta.github.io/ESPNet/>.

1 Introduction

Deep convolutional neural network (CNN) models have achieved high accuracy in visual scene understanding tasks [1–3]. While the accuracy of these networks has improved with their increase in depth and width, large networks are slow and power hungry. This is especially problematic on the computationally heavy task of semantic segmentation [4–10]. For example, PSPNet [1] has 65.7 million parameters and runs at about 1 FPS while discharging the battery of a standard laptop at a rate of 77 Watts. Many advanced real-world applications, such as self-driving cars, robots, and augmented reality, are sensitive and demand on-line

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-01249-6_34) contains supplementary material, which is available to authorized users.

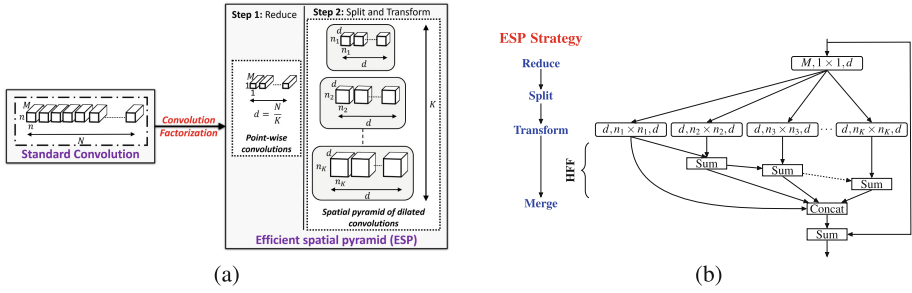


Fig. 1. (a) The standard convolution layer is decomposed into point-wise convolution and spatial pyramid of dilated convolutions to build an efficient spatial pyramid (ESP) module. (b) Block diagram of ESP module. The large effective receptive field of the ESP module introduces gridding artifacts, which are removed using hierarchical feature fusion (HFF). A skip-connection between input and output is added to improve the information flow. See Sect. 3 for more details. Dilated convolutional layers are denoted as $(\# \text{ input channels, effective kernel size, } \# \text{ output channels})$. The effective spatial dimensions of a dilated convolutional kernel are $n_k \times n_k$, where $n_k = (n-1)2^{k-1} + 1, k = 1, \dots, K$. Note that only $n \times n$ pixels participate in the dilated convolutional kernel. In our experiments $n = 3$ and $d = \frac{M}{K}$.

processing of data locally on edge devices. These accurate networks require enormous resources and are not suitable for edge devices, which have limited energy overhead, restrictive memory constraints, and reduced computational capabilities.

Convolution factorization has demonstrated its success in reducing the computational complexity of deep CNNs [11–15]. We introduce an efficient convolutional module, ESP (efficient spatial pyramid), which is based on the convolutional factorization principle (Fig. 1). Based on these ESP modules, we introduce an efficient network structure, ESPNet, that can be easily deployed on resource-constrained edge devices. ESPNet is *fast, small, low power, and low latency*, yet still preserves segmentation accuracy.

ESP is based on a convolution factorization principle that decomposes a standard convolution into two steps: (1) *point-wise convolutions* and (2) *spatial pyramid of dilated convolutions*, as shown in Fig. 1. The point-wise convolutions help in reducing the computation, while the spatial pyramid of dilated convolutions re-samples the feature maps to learn the representations from large effective receptive field. We show that our ESP module is more efficient than other factorized forms of convolutions, such as Inception [11–13] and ResNext [14]. Under the same constraints on memory and computation, ESPNet outperforms MobileNet [16] and ShuffleNet [17] (two other efficient networks that are built upon the factorization principle). We note that existing spatial pyramid methods (e.g. the atrous spatial pyramid module in [3]) are computationally expensive and cannot be used at different spatial levels for learning the representations. In contrast to these methods, ESP is computationally efficient and

can be used at different spatial levels of a CNN network. Existing models based on dilated convolutions [1, 3, 18, 19] are large and inefficient, but our ESP module generalizes the use of dilated convolutions in a novel and efficient way.

To analyze the performance of a CNN network on edge devices, we introduce several new performance metrics, such as sensitivity to GPU frequency and warp execution efficiency. To showcase the power of ESPNet, we evaluate our model on one of the most expensive tasks in AI and computer vision: semantic segmentation. ESPNet is empirically demonstrated to be more accurate, efficient, and fast than ENet [20], one of the most power-efficient semantic segmentation networks, while learning a similar number of parameters. Our results also show that ESPNet learns generalizable representations and outperforms ENet [20] and another efficient network ERFNet [21] on the unseen dataset. ESPNet can process a *high resolution RGB image* at a rate of 112, 21, and 9 frames per second on the NVIDIA TitanX, GTX-960M, and Jetson TX2 respectively.

2 Related Work

Different techniques, such as convolution factorization, network compression, and low-bit networks, have been proposed to speed up CNNs. We, first, briefly describe these approaches and then provide a brief overview of CNN-based semantic segmentation.

Convolution Factorization: Convolutional factorization decomposes the convolutional operation into multiple steps to reduce the computational complexity. This factorization has successfully shown its potential in reducing the computational complexity of deep CNN networks (e.g. Inception [11–13], factorized network [22], ResNext [14], Xception [15], and MobileNets [16]). ESP modules are also built on this factorization principle. The ESP module decomposes a convolutional layer into a point-wise convolution and spatial pyramid of dilated convolutions. This factorization helps in reducing the computational complexity, while simultaneously allowing the network to learn the representations from a large effective receptive field. **Network Compression:** Another approach for building efficient networks is compression. These methods use techniques such as hashing [23], pruning [24], vector quantization [25], and shrinking [26, 27] to reduce the size of the pre-trained network. **Low-bit networks:** Another approach towards efficient networks is low-bit networks, which quantize the weights to reduce the network size and complexity (e.g. [28–31]). **Sparse CNN:** To remove the redundancy in CNNs, sparse CNN methods, such as sparse decomposition [32], structural sparsity learning [33], and dictionary-based method [34], have been proposed.

We note that compression-based methods, low-bit networks, and sparse CNN methods are equally applicable to ESPNets and are complementary to our work.

Dilated Convolution: Dilated convolutions [35] are a special form of standard convolutions in which the effective receptive field of kernels is increased by inserting zeros (or holes) between each pixel in the convolutional kernel. For a

$n \times n$ dilated convolutional kernel with a dilation rate of r , the effective size of the kernel is $[(n-1)r+1]^2$. The dilation rate specifies the number of zeros (or holes) between pixels. However, due to dilation, only $n \times n$ pixels participate in the convolutional operation, reducing the computational cost while increasing the effective kernel size.

Yu and Koltun [18] stacked dilated convolution layers with increasing dilation rate to learn contextual representations from a large effective receptive field. A similar strategy was adopted in [19, 36, 37]. Chen et al. [3] introduced an atrous spatial pyramid (ASP) module. This module can be viewed as a parallelized version of [3]. These modules are computationally inefficient (e.g. ASPs have high memory requirements and learn many more parameters; see Sect. 3.2). Our ESP module also learns multi-scale representations using dilated convolutions in parallel; however, it is computationally efficient and can be used at any spatial level of a CNN network.

CNN for Semantic Segmentation: Different CNN-based segmentation networks have been proposed, such as multi-dimensional recurrent neural networks [38], encoder-decoders [20, 21, 39, 40], hypercolumns [41], region-based representations [42, 43], and cascaded networks [44]. Several supporting techniques along with these networks have been used for achieving high accuracy, including ensembling features [3], multi-stage training [45], additional training data from other datasets [1, 3], object proposals [46], CRF-based post processing [3], and pyramid-based feature re-sampling [1–3].

Encoder-Decoder Networks: Our work is related to this line of work. The encoder-decoder networks first learn the representations by performing convolutional and down-sampling operations. These representations are then decoded by performing up-sampling and convolutional operations. ESPNet first learns the encoder and then attaches a *light-weight decoder* to produce the segmentation mask. This is in contrast to existing networks where the decoder is either an exact replica of the encoder (e.g. [39]) or is relatively small (but not light weight) in comparison to the encoder (e.g. [20, 21]).

Feature Re-sampling Methods: The feature re-sampling methods re-sample the convolutional feature maps at the same scale using different pooling rates [1, 2] and kernel sizes [3] for efficient classification. Feature re-sampling is computationally expensive and is performed just before the classification layer to learn scale-invariant representations. We introduce a computationally efficient convolutional module that allows feature re-sampling at different spatial levels of a CNN network.

3 ESPNet

We describe ESPNet and its core ESP module. We compare ESP modules with similar CNN modules, Inception [11–13], ResNext [14], MobileNet [16], and ShuffleNet [17].

3.1 ESP Module

ESPNet is based on efficient spatial pyramid (ESP) modules, a factorized form of convolutions that decompose a standard convolution into a point-wise convolution and a spatial pyramid of dilated convolutions (see Fig. 1a). The point-wise convolution applies a 1×1 convolution to project high-dimensional feature maps onto a low-dimensional space. The spatial pyramid of dilated convolutions then re-samples these low-dimensional feature maps using K , $n \times n$ dilated convolutional kernels simultaneously, each with a dilation rate of 2^{k-1} , $k = \{1, \dots, K\}$. This factorization drastically reduces the number of parameters and the memory required by the ESP module, while preserving a large effective receptive field $[(n-1)2^{K-1} + 1]^2$. This pyramidal convolutional operation is called a spatial pyramid of dilated convolutions, because each dilated convolutional kernel learns weights with different receptive fields and so resembles a spatial pyramid.

A standard convolutional layer takes an input feature map $\mathbf{F}_i \in \mathbb{R}^{W \times H \times M}$ and applies N kernels $\mathbf{K} \in \mathbb{R}^{m \times n \times M}$ to produce an output feature map $\mathbf{F}_o \in \mathbb{R}^{W \times H \times N}$, where W and H represent the width and height of the feature map, m and n represent the width and height of the kernel, and M and N represent the number of input and output feature channels. For simplicity, we will assume that $m = n$. A standard convolutional kernel thus learns $n^2 MN$ parameters. These parameters are *multiplicatively* dependent on the spatial dimensions of the $n \times n$ kernel and the number of input M and output N channels.

Width Divider K : To reduce the computational cost, we introduce a simple hyper-parameter K . The role of K is to shrink the dimensionality of the feature maps uniformly across each ESP module in the network. *Reduce:* For a given K , the ESP module first reduces the feature maps from M -dimensional space to $\frac{N}{K}$ -dimensional space using a point-wise convolution (Step 1 in Fig. 1a). *Split:* The low-dimensional feature maps are split across K parallel branches. *Transform:* Each branch then processes these feature maps simultaneously using $n \times n$ dilated convolutional kernels with different dilation rates given by 2^{k-1} , $k = \{1, \dots, K-1\}$ (Step 2 in Fig. 1a). *Merge:* The outputs of the K parallel dilated convolutional kernels are concatenated to produce an N -dimensional output feature map Fig. 1b visualizes the *reduce-split-transform-merge* strategy.

The ESP module has $(NM + (Nn)^2)/K$ parameters and its effective receptive field is $((n-1)2^{K-1} + 1)^2$. Compared to the $n^2 NM$ parameters of the standard convolution, factorizing it reduces the number of parameters by a factor of $\frac{n^2 MK}{M+n^2 N}$, while increasing the effective receptive field by $\sim (2^{K-1})^2$. For example, the ESP module learns $\sim 3.6 \times$ fewer parameters with an effective receptive field of 17×17 than a standard convolutional kernel with an effective receptive field of 3×3 for $n = 3$, $N = M = 128$, and $K = 4$.

Hierarchical Feature Fusion (HFF) for De-gridding: While concatenating the outputs of dilated convolutions give the ESP module a large effective receptive field, it introduces unwanted checkerboard or gridding artifacts, as shown in Fig. 2. To address the gridding artifact in ESP, the feature maps obtained using

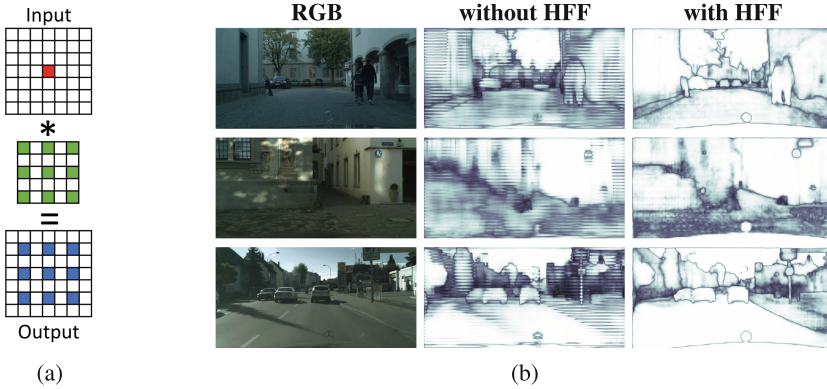


Fig. 2. (a) An example illustrating a gridding artifact with a single active pixel (red) convolved with a 3×3 dilated convolutional kernel with dilation rate $r = 2$. (b) Visualization of feature maps of ESP modules with and without hierarchical feature fusion (HFF). HFF in ESP eliminates the gridding artifact. Best viewed in color.

kernels of different dilation rates are hierarchically added before concatenating them (HFF in Fig. 1b). This simple, effective solution does not increase the complexity of the ESP module, in contrast to existing methods that remove the gridding artifact by learning more parameters using dilated convolutional kernels [19, 37]. To improve gradient flow inside the network, the input and output feature maps are combined using an element-wise sum [47].

3.2 Relationship with Other CNN Modules

The ESP module shares similarities with the following CNN modules.

MobileNet Module: The MobileNet module [16], shown in Fig. 3a, uses a depth-wise separable convolution [15] that factorizes a standard convolutions into depth-wise convolutions (*transform*) and point-wise convolutions (*expand*). It learns less parameters, has high memory requirement, and low receptive field than the ESP module. An extreme version of the ESP module (with $K = N$) is almost identical to the MobileNet module, differing only in the order of convolutional operations. In the MobileNet module, the spatial convolutions are followed by point-wise convolutions; however, in the ESP module, point-wise convolutions are followed by spatial convolutions.

ShuffleNet Module: The ShuffleNet module [17], shown in Fig. 3b, is based on the principle of *reduce-transform-expand*. It is an optimized version of the bottleneck block in ResNet [47]. To reduce computation, Shufflenet makes use of grouped convolutions [48] and depth-wise convolutions [15]. It replaces 1×1 and 3×3 convolutions in the bottleneck block in ResNet with 1×1 grouped convolutions and 3×3 depth-wise separable convolutions, respectively. The Shufflenet module learns many less parameters than the ESP module, but has higher memory requirements and a smaller receptive field.

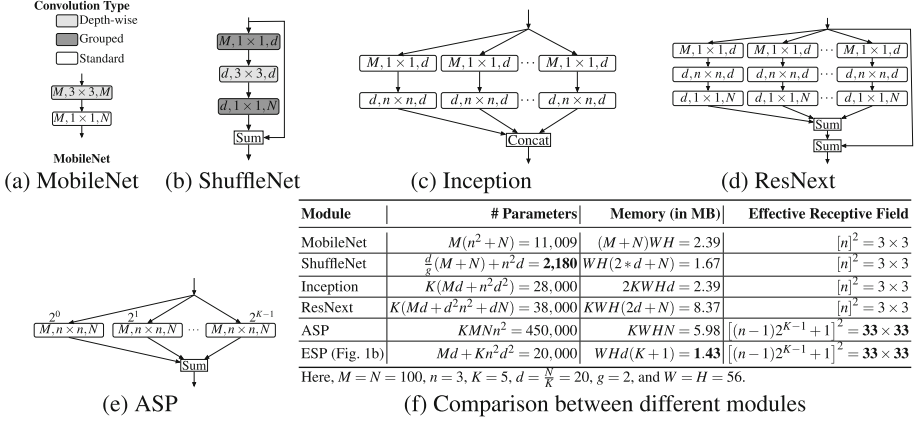


Fig. 3. Different types of convolutional modules for comparison. We denote the layer as (# input channels, kernel size, # output channels). Dilation rate in (e) is indicated on top of each layer. Here, g represents the number of convolutional groups in grouped convolution [48]. For simplicity, we only report the memory of convolutional layers in (d). For converting the required memory to bytes, we multiply it by 4 (1 float requires 4 bytes for storage).

Inception Module: Inception modules [11–13] are built on the principle of *split-reduce-transform-merge* and are usually heterogeneous in number of channels and kernel size (e.g. some of the modules are composed of standard and factored convolutions). In contrast, ESP modules are straightforward and simple to design. For the sake of comparison, the homogeneous version of an Inception module is shown in Fig. 3c. Figure 3f compares the Inception module with the ESP module. ESP (1) learns fewer parameters, (2) has a low memory requirement, and (3) has a larger effective receptive field.

ResNext Module: A ResNext module [14], shown in Fig. 3d, is a parallel version of the bottleneck module in ResNet [47], based on the principle of *split-reduce-transform-expand-merge*. The ESP module is similar in branching and residual summation, but more efficient in memory and parameters with a larger effective receptive field.

Atrous Spatial Pyramid (ASP) Module: An ASP module [3], shown in Fig. 3e, is built on the principle of *split-transform-merge*. The ASP module involves branching with each branch learning kernel at a different receptive field (using dilated convolutions). Though ASP modules tend to perform well in segmentation tasks due to their high effective receptive fields, ASP modules have high memory requirements and learn many more parameters. Unlike the ASP module, the ESP module is computationally efficient.

4 Experiments

To showcase the power of ESPNet, we evaluate ESPNet’s performance on several semantic segmentation datasets and compare to the state-of-the-art networks.

4.1 Experimental Set-Up

Network Structure: ESPNet uses ESP modules for learning convolutional kernels as well as down-sampling operations, except for the first layer: a standard strided convolution. All layers are followed by a batch normalization [49] and a PReLU [50] non-linearity except the last point-wise convolution, which has neither batch normalization nor non-linearity. The last layer feeds into a softmax for pixel-wise classification.

Different variants of ESPNet are shown in Fig. 4. The first variant, ESPNet-A (Fig. 4a), is a standard network that takes an RGB image as an input and learns representations at different spatial levels¹ using the ESP module to produce a segmentation mask. The second variant, ESPNet-B (Fig. 4b), improves the flow of information inside ESPNet-A by sharing the feature maps between the previous strided ESP module and the previous ESP module. The third variant, ESPNet-C (Fig. 4c), reinforces the input image inside ESPNet-B to further improve the flow of information. These three variants produce outputs whose spatial dimensions are $\frac{1}{8}th$ of the input image. The fourth variant, ESPNet (Fig. 4d), adds a light weight decoder (built using a principle of *reduce-upsample-merge*) to ESPNet-C that outputs the segmentation mask of the same spatial resolution as the input image.

To build deeper computationally efficient networks for edge devices without changing the network topology, a hyper-parameter α controls the depth of the network; the ESP module is repeated α_l times at spatial level l . CNNs require more memory at higher spatial levels (at $l = 0$ and $l = 1$) because of the high spatial dimensions of feature maps at these levels. To be memory efficient, neither the ESP nor the convolutional modules are repeated at these spatial levels.

Dataset: We evaluated the ESPNet on the Cityscapes dataset [6], an urban visual scene-understanding dataset that consists of 2,975 training, 500 validation, and 1,525 test high-resolution images. The task is to segment an image into 19 classes belonging to 7 categories (e.g. person and rider classes belong to the same category *human*). We evaluated our networks on the test set using the Cityscapes *online* server.

To study the generalizability, we tested the ESPNet on an unseen dataset. We used the Mapillary dataset [51] for this task because of its diversity. We mapped the annotations (65 classes) in the validation set (# 2,000 images) to seven categories in the Cityscape dataset. To further study the segmentation power of our model, we trained and tested the ESPNet on two other popular

¹ At each spatial level l , the spatial dimensions of the feature maps are the same. To learn representations at different spatial levels, a down-sampling operation is performed (see Fig. 4a).

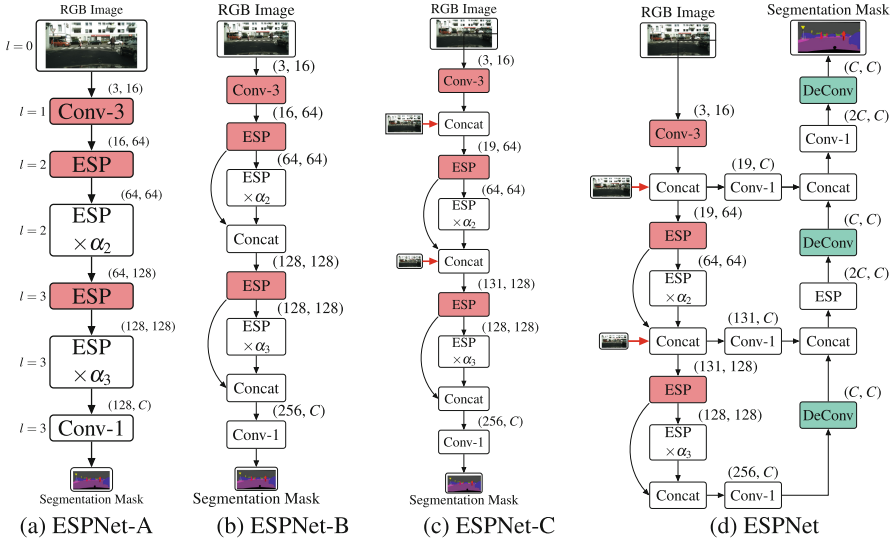


Fig. 4. The path from ESPNet-A to ESPNet. Red and green color boxes represent the modules responsible for down-sampling and up-sampling operations, respectively. Spatial-level l is indicated on the left of every module in (a). We denote each module as $(\# \text{ input channels}, \# \text{ output channels})$. Here, Conv- n represents $n \times n$ convolution. (Color figure online)

datasets from different domains. First, we used the widely known PASCAL VOC dataset [52] that has 1,464 training images, 1,448 validation images, and 1,456 test images. The task is to segment an image into 20 foreground classes. We evaluate our networks on the test set (comp6 category) using the PASCAL VOC *online* server. Following the convention, we used additional images from [53, 54]. Secondly, we used a breast biopsy whole slide image dataset [36], chosen because tissue structures in biomedical images vary in size and shape and because this dataset allowed us to check the potential of learning representations from a large receptive field. The dataset consists of 30 training images and 28 validation images, whose average size is $10,000 \times 12,000$ pixels, much larger than natural scene images. The task is to segment the images into 8 biological tissue labels; details are in [36].

Performance Evaluation Metrics: Most traditional CNNs measure network performance in terms of accuracy, latency, network parameters, and network size [16, 17, 20, 21, 55]. These metrics provide high-level insight about the network, but fail to demonstrate the efficient usage of hardware resources with limited availability. In addition to these metrics, we introduce several *system-level metrics* to characterize the performance of a CNN on resource-constrained devices [56, 57].

Segmentation accuracy is measured as a mean Intersection over Union (mIOU) score between the ground truth and the predicted segmentation mask.

Latency represents the amount of time a CNN network takes to process an image. This is usually measured in terms of frames per second (FPS).

Network parameters represents the number of parameters learned by the network.

Network size represents the amount of storage space required to store the network parameters. An efficient network should have a smaller network size.

Power consumption is the average power consumed by the network during inference.

Sensitivity to GPU frequency measures the computational capability of an application and is defined as a ratio of percentage change in execution time to the percentage change in GPU frequency. Higher values indicate higher efficiency.

Utilization rates measure the utilization of compute resources (CPU, GPU, and memory) while running on an edge device. In particular, computing units in edge devices (e.g. Jetson TX2) share memory between CPU and GPU.

Warp execution efficiency is defined as the average percentage of active threads in each executed warp. GPUs schedule threads as warps; each thread is executed in *single instruction multiple data* fashion. Higher values represent efficient usage of GPU.

Memory efficiency is the ratio of number of bytes requested/stored to the number of bytes transferred from/to device (or shared) memory to satisfy load/store requests. Since memory transactions are in blocks, this metric measures memory bandwidth efficiency.

Training Details: ESPNet networks were trained using PyTorch [58] with CUDA 9.0 and cuDNN back-ends. ADAM [59] was used with an initial learning rate of 0.0005, and decayed by two after every 100 epochs and with a weight decay of 0.0005. An inverse class probability weighting scheme was used in the cross-entropy loss function to address the class imbalance [20,21]. Following [20,21], the weights were initialized randomly. Standard strategies, such as scaling, cropping and flipping, were used to augment the data. The image resolution in the Cityscape dataset is 2048×1024 , and all the accuracy results were reported at this resolution. For training the networks, we sub-sampled the RGB images by two. When the output resolution was smaller than 2048×1024 , the output was up-sampled using bi-linear interpolation. For training on the PASCAL dataset, we used a fixed image size of 512×512 . For the WSI dataset, the patch-wise training approach was followed [36]. ESPNet was trained in two stages. First, ESPNet-C was trained with down-sampled annotations. Second, a light-weight decoder was attached to ESPNet-C and then, the entire ESPNet network was trained.

Three different GPU devices were used for our experiments: (1) a desktop with a NVIDIA TitanX GPU (3,584 CUDA cores), (2) a laptop with a NVIDIA GTX-960M GPU (640 CUDA cores), and (3) an edge device with a NVIDIA Jetson TX2 (256 CUDA cores). Unless and otherwise stated explicitly, statistics

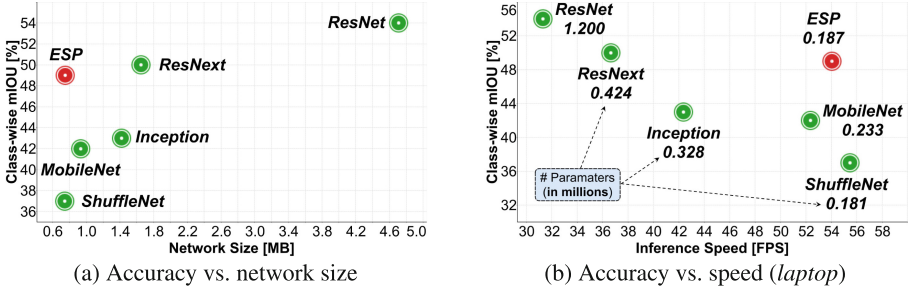


Fig. 5. Comparison between state-of-the-art efficient convolutional modules. For a fair comparison between different modules, we used $K = 5$, $d = \frac{N}{K}$, $\alpha_2 = 2$, and $\alpha_3 = 3$. We used standard strided convolution for down-sampling. For ShuffleNet, we used $g = 4$ and $K = 4$ so that the resultant ESPNet-C network has the same complexity as with the ESP block.

are reported for an RGB image of size 1024×512 averaged over 200 trials. For collecting the hardware-level statistics, NVIDIA’s and Intel’s hardware profiling and tracing tools, such as NVPROF [60], Tegrastats [61], and PowerTop [62], were used. In our experiments, we will refer to ESPNet with $\alpha_2 = 2$ and $\alpha_3 = 8$ as ESPNet until and otherwise stated explicitly.

4.2 Segmentation Results on the Cityscape Dataset

Comparison with Efficient Convolutional Modules: In order to understand the ESP module, we replaced the ESP modules in ESPNet-C with state-of-the-art efficient convolutional modules, sketched in Fig. 3 (MobileNet [16], ShuffleNet [17], Inception [11–13], ResNext [14], and ResNet [47]) and evaluated their performance on the Cityscape validation dataset. We did not compare with ASP [3], because it is computationally expensive and not suitable for edge devices. Figure 5 compares the performance of ESPNet-C with different convolutional modules. Our ESP module outperformed MobileNet and ShuffleNet modules by 7% and 12%, respectively, while learning a similar number of parameters and having comparable network size and inference speed. Furthermore, the ESP module delivered comparable accuracy to ResNext and Inception more efficiently. A basic ResNet module (stack of two 3×3 convolutions with a skip-connection) delivered the best performance, but had to learn $6.5\times$ more parameters.


Comparison with Segmentation Methods: We compared the performance of ESPNet with state-of-the-art semantic segmentation networks. These networks either use a pre-trained network (VGG [63]: FCN-8s [45] and SegNet [39], ResNet [47]: DeepLab-v2 [3] and PSPNet [1], and SqueezeNet [55]: SQNet [64]) or were trained from scratch (ENet [20] and ERFNet [21]). ESPNet is 2% more accurate than ENet [20], while running $1.27\times$ and $1.16\times$ faster on a desktop and a laptop, respectively (Fig. 6). ESPNet makes some mistakes between classes

that belong to the same category, and hence has a lower class-wise accuracy. For example, a rider can be confused with a person. However, ESPNet delivers a good category-wise accuracy. ESPNet had 8% lower category-wise mIOU than PSPNet [1], while learning 180× fewer parameters. ESPNet had lower power consumption, had lower battery discharge rate, and was significantly faster than state-of-the-art methods, while still achieving a competitive category-wise accuracy; this makes ESPNet suitable for segmentation on edge devices. ERFNet, an another efficient segmentation network, delivered good segmentation accuracy, but has 5.5× more parameters, is 5.44× larger, consumes more power, and has a higher battery discharge rate than ESPNet. Also, ERFNet does not utilize limited available hardware resources efficiently on edge devices (Sect. 4.4).


| Network | mIOU | |
|---------------|-------|----------|
| | Class | Category |
| ENet [20] | 58.3 | 80.4 |
| ERFNet [21] | 68.0 | 86.5 |
| SQNet [27] | 59.8 | 84.3 |
| SegNet [39] | 57.0 | 79.1 |
| ESPNet (Ours) | 60.3 | 82.2 |

| | | |
|----------------|------|------|
| FCN-8s [39] | 65.3 | 85.7 |
| DeepLab-v2 [3] | 70.4 | 86.4 |
| PSPNet [1] | 78.4 | 90.6 |


(a) Test set




(b) Accuracy vs. network size




(c) Accuracy vs. # parameters




(d) Battery discharge rate vs. network (*laptop*)



(e) Accuracy vs. speed (*laptop*)



(f) Power consumption vs. speed (*laptop*)



(g) Power consumption vs. speed (*desktop*)

Fig. 6. Comparison between segmentation methods on the Cityscape test set on two different devices. All networks (FCN-8s [45], SegNet [39], SQNet [64], ENet [20], DeepLab-v2 [3], PSPNet [1], and ERFNet [21]) were without CRF and converted to PyTorch for a fair comparison.

4.3 Segmentation Results on Other Datasets

Unseen Dataset: Table 1a compares the performance of ESPNet with ENet [20] and ERFNet [21] on an unseen dataset. These networks were trained on the Cityscapes dataset [6] and tested on the Mapillary (unseen) dataset [51]. ENet and ERFNet were chosen, due to the efficiency and power of ENet and high accuracy of ERFNet. Our experiments show that ESPNet learns good generalizable representations of objects and outperforms ENet and ERFNet on the unseen dataset.


PASCAL VOC 2012 Dataset: (Table 1c) On the PASCAL dataset, ESPNet is 4% more accurate than SegNet, one of the smallest network on the PASCAL VOC, while learning $81\times$ fewer parameters. ESPNet is 22% less accurate than PSPNet (one of the most accurate network on the PASCAL VOC) while learning $180\times$ fewer parameters.

Breast Biopsy Dataset: (Table 1d) On the breast biopsy dataset, ESPNet achieved the same accuracy as [36] while learning $9.5\times$ less parameters.

Table 1. Results on different datasets, where $^\circ$ denotes the values are in millions. * See [66].

| | mIOU | # Params $^\circ$ |
|-------------|-------------|-------------------|
| ENet [20] | 0.33 | 0.364 |
| ERFNet [21] | 0.25 | 2.06 |
| ESPNet | 0.40 | 0.364 |

(a) Mapillary validation set [51]



(b) Mapillary validation set [51] (unseen)

| Model | ESPNet (Ours) | SegNet [39] | RefineNet [44] | DeepLab [3] | PSPNet [1] | LRR [65] | Dilation-8 [18] | FCN-8s [45] |
|-------------------|---------------|-------------|----------------|-------------|--------------|----------|-----------------|-------------|
| # Params $^\circ$ | 0.364 | 29.5 | 42.6 | 44.04 | 65.7 | 48 | 141.13 | 134.5 |
| mIOU | 63.01 | 59.10 | 82.40 | 79.70 | 85.40 | 79.30 | 75.30 | 67.20 |

(c) PASCAL VOC test set [52]

| Model | Module | mIOU | # Params $^\circ$ |
|--------------------------|--------|--------------|-------------------|
| SegNet [39] | VGG | 37.6 | 12.80 |
| Mehta <i>et al.</i> [36] | ResNet | 44.20 | 26.03 |
| ESPNet* | ESP | 44.03 | 2.75 |

(d) Breast biopsy validation set [36]

4.4 Performance Analysis on the NVIDIA Jetson TX2 (Edge Device)

Network Size: Figure 7a compares the uncompressed 32-bit network size of ESPNet with ENet and ERFNet. ESPNet had a $1.12\times$ and $5.45\times$ smaller network than ENet and ERFNet, respectively, which reflects well on the architectural design of ESPNet.

Inference Speed and Sensitivity to GPU Frequency: Figure 7b compares the inference speed of ESPNet with ENet and ERFNet. ESPNet had almost the same frame rate as ENet, but it was more sensitive to GPU frequency (Fig. 7c).

As a consequence, ESPNet achieved a higher frame rate than ENet on high-end graphic cards, such as the GTX-960M and TitanX (see Fig. 6). For example, ESPNet is 1.27× faster than ENet on an NVIDIA TitanX. ESPNet is about 3× faster than ERFNet on an NVIDIA Jetson TX2.

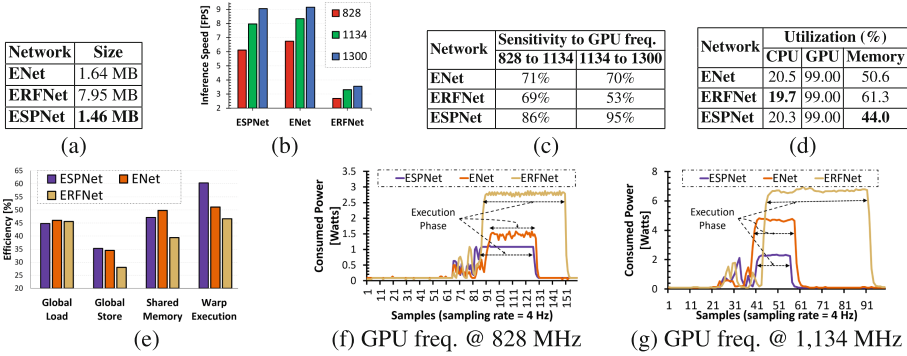


Fig. 7. Performance analysis of ESPNet with ENet and ERFNet on a NVIDIA Jetson TX2: (a) network size, (b) inference speed vs. GPU frequency (in MHz), (c) sensitivity analysis, (d) utilization rates, (e) efficiency rates, and (f, g) power consumption at two different GPU frequencies. In (d), initialization phase statistics were not considered, due to similarity across all networks.

Utilization Rates: Figure 7d compares the CPU, GPU, and memory utilization rates of networks that are throughput intensive; GPU utilization rates are high, while CPU utilization rates are low for these networks. Memory utilization rates are significantly different for these networks. The memory footprint of ESPNet is low in comparison to ENet and ERFNet, suggesting that ESPNet is suitable for memory-constrained devices.

Warp Execution Efficiency: Figure 7e compares the warp execution efficiency of ESPNet with ENet and ERFNet. The warp execution of ESPNet was about 9% higher than ENet and about 14% higher than ERFNet. This indicates that ESPNet has less warp divergence and promotes the efficient usage of limited GPU resources available on edge devices. We note that warp execution efficiency gives a better insight into the utilization of GPU resources than the GPU utilization rate. GPU frequency will be busy even if few warps are active, resulting in a high GPU utilization rate.

Memory Efficiency: (Figure 7e) All networks have similar global load efficiency, but ERFNet has a poor store and shared memory efficiency. This is likely due to the fact that ERFNet spends 20% of the compute power performing memory alignment operations, while ESPNet and ENet spend 4.2% and 6.6% time for this operation, respectively.

Power Consumption: Figure 7f and g compares the power consumption of ESPNet with ENet and ERFNet at two different GPU frequencies. The average power consumption (during network execution phase) of ESPNet, ENet, and ERFNet were 1 W, 1.5 W, and 2.9 W at a GPU frequency of 824 MHz and 2.2 W, 4.6 W, and 6.7 W at a GPU frequency of 1,134 MHz, respectively; suggesting ESPNet is a power-efficient network.

4.5 Ablation Studies on the Cityscapes: The Path from ESPNet-A to ESPNet

Larger networks or ensembling the output of multiple networks delivers better performance [1, 3, 19], but with ESPNet (sketched in Fig. 4), the goal is an efficient network for edge devices. To improve the performance of ESPNet while maintaining efficiency, a systematic study of design choices was performed. Table 2 summarizes the results.

Table 2. The path from ESPNet-A to ESPNet. Here, ERF represents effective receptive field, * denotes that strided ESP was used for down-sampling, † indicates that the input reinforcement method was replaced with input-aware fusion method [36], and ° denotes the values are in million. All networks in (a–c, e–f) are trained for 100 epochs, while networks in (d, g) are trained for 300 epochs. Here, SPC-s denotes that 3×3 standard convolutions are used instead of dilated convolutions in the spatial pyramid of dilated convolutions (SPC).

| mIOU # Params° | | | Module mIOU # Params° | | | Downsample mIOU # Params° | | | ESPNet-C configuration | | | | | |
|----------------|-------------|-------|------------------------|-------------|-------|---------------------------|-------------|--------------|-------------------------|--------------|-------|-------|--------|------|
| ReLU | 0.36 | 0.183 | ESP | 0.39 | 0.183 | Strided conv. | 0.38 | 0.274 | C1 - ($\alpha_3 = 3$) | 3×3 | SPC | 0.276 | 1.2 MB | 50.8 |
| PReLU | 0.38 | 0.183 | -RL | 0.37 | 0.183 | Strided ESP | 0.39 | 0.183 | C2 - ($\alpha_3 = 3$) | 1×1 | SPC | 0.187 | 0.8 MB | 49.0 |
| | | | RL - residual learning | | | | | | C3 - ($\alpha_3 = 3$) | 1×1 | SPC-s | 0.187 | 0.8 MB | 47.4 |

(a)
(b)
(c)
(d)

| mIOU # Params° | Width divider K | | | | | | Network mIOU # Params° | | | α_3 mIOU # Params (in million) Network size | | | ESPNet (Fig. 4d) mIOU # Params (in million) Network size | | | |
|----------------------------|-------------------|-----------------|-----------------|-----------------|------------------|------------------------|-------------------------|-------------|--------------|--|-------------|--------------|--|-------------|--------------|----------------|
| | 2 | 4 | 5 | 6 | 7 | 8 | Network | mIOU | # Params° | α_3 | mIOU | # Params | Network size | mIOU | # Params | Network size |
| mIOU | 0.415 | 0.378 | 0.381 | 0.359 | 0.321 | 0.303 | ESPNet-A* | 0.39 | 0.183 | 3 | 49.0 | 0.187 | 0.75 MB | 56.3 | 0.202 | 0.82 MB |
| # Params° | 0.358 | 0.215 | 0.183 | 0.165 | 0.152 | 0.143 | ESPNet-B | 0.40 | 0.186 | 5 | 51.2 | 0.252 | 1.01 MB | 57.9 | 0.267 | 1.07 MB |
| ERF ($n^2 = n \times n$) | 5 ² | 17 ² | 33 ² | 65 ² | 129 ² | 257² | ESPNet-C | 0.42 | 0.187 | 8 | 53.3 | 0.349 | 1.40 MB | 61.4 | 0.364 | 1.46 MB |
| | | | | | | | ESPNet-C† | 0.42 | 0.206 | | | | | | | |

(e)
(f)
(g)

ReLU vs PReLU: (Table 2a) Replacing ReLU [67] with PReLU [50] in ESPNet-A improved the accuracy by 2%, while having a minimal impact on the network complexity.

Residual Learning in ESP: (Table 2b) The accuracy of ESPNet-A dropped by about 2% when skip-connections in ESP (Fig. 1b) modules were removed. This verifies the effectiveness of the residual learning.

Down-Sampling: (Table 2c) Replacing the standard strided convolution with the strided ESP in ESPNet-A improved accuracy by 1% with 33% parameter reduction.

Width Divider (K): (Table 2e) Increasing K enlarges the effective receptive field of the ESP module, while simultaneously decreasing the number of network parameters. Importantly, ESPNet-A’s accuracy decreased with increasing K . For example, raising K from 2 to 8 caused ESPNet-A’s accuracy to drop by 11%. This drop in accuracy is explained in part by the ESP module’s effective receptive field growing beyond the size of its input feature maps. For an image with size 1024×512 , the spatial dimensions of the input feature maps at spatial level $l = 2$ and $l = 3$ are 256×128 and 128×64 , respectively. However, some of the kernels have larger receptive fields (257×257 for $K = 8$). The weights of such kernels do not contribute to learning, thus resulting in lower accuracy. At $K = 5$, we found a good trade-off between number of parameters and accuracy, and therefore, we used $K = 5$ in our experiments.

ESPNet-A \rightarrow ESPNet-C: (Table 2f) Replacing the convolution-based network width expansion operation in ESPNet-A with the concatenation operation in ESPNet-B improved the accuracy by about 1% and did not increase the number of network parameters noticeably. With input reinforcement (ESPNet-C), the accuracy of ESPNet-B further improved by about 2%, while not increasing the network parameters drastically. This is likely due to the fact that the input reinforcement method establishes a direct link between the input image and encoding stage, improving the flow of information.

The closest work to our input reinforcement method is the input-aware fusion method of [36], which learns representations on the down-sampled input image and additively combines them with the convolutional unit. When the proposed input reinforcement method was replaced with the input-aware fusion in [36], no improvement in accuracy was observed, but the number of network parameters increased by about 10%.

ESPNet-C vs ESPNet: (Table 2g) Adding a light-weight decoder to ESPNet-C improved the accuracy by about 6%, while increasing the number of parameters and network size by merely 20,000 and 0.06 MB from ESPNet-C to ESPNet, respectively.

Impact of Different Convolutions in the ESP Block: The ESP block uses point-wise convolutions for reducing the high-dimensional feature maps to low-dimensional space and then transforms those feature maps using a spatial pyramid of dilated convolutions (SPCs) (see Sect. 3). To understand the influence of these two components, we performed the following experiments. (1) *Point-wise convolutions:* We replaced point-wise convolutions with 3×3 standard convolutions in the ESP block (see C1 and C2 in Table 2d), and the resultant network demanded more resources (e.g., 47% more parameters) while improving the mIOU by 1.8%, showing that point-wise convolutions are effective. Moreover, the decrease in number of parameters due to point-wise convolutions in the ESP block enables the construction of deep and efficient networks (see Table 2g). (2) *SPCs:* We replaced 3×3 dilated convolutions with 3×3 standard convolutions in the ESP block. Though the resultant network is as efficient as with dilated

convolutions, it is 1.6% less accurate; suggesting SPCs are effective (see C2 and C3 in Table 2d).

5 Conclusion

We introduced a semantic segmentation network, ESPNet, based on an efficient spatial pyramid module. In addition to legacy metrics, we introduced several new system-level metrics that help to analyze the performance of a CNN network. Our empirical analysis suggests that ESPNets are fast and efficient. We also demonstrated that ESPNet learns good generalizable representations of the objects and perform well in the wild.

Acknowledgement. This research was supported by the Intelligence Advanced Research Projects Activity (IARPA) via Interior/Interior Business Center (DOI/IBC) contract number D17PC00343, the Washington State Department of Transportation research grant T1461-47, NSF III (1703166), the National Cancer Institute awards (R01 CA172343, R01 CA140560, and R01 CA200690), Allen Distinguished Investigator Award, Samsung GRO award, and gifts from Google, Amazon, and Bloomberg. We would also like to acknowledge NVIDIA Corporation for donating the Jetson TX2 board and the Titan X Pascal GPU used for this research. We also thank the anonymous reviewers for their helpful comments. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing endorsements, either expressed or implied, of IARPA, DOI/IBC, or the U.S. Government.

References

1. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: CVPR (2017)
2. He, K., Zhang, X., Ren, S., Sun, J.: Spatial pyramid pooling in deep convolutional networks for visual recognition. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8691, pp. 346–361. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10578-9_23
3. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: DeepLab: semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. TPAMI **40**, 834–848 (2018)
4. Ess, A., Müller, T., Grabner, H., Van Gool, L.J.: Segmentation-based urban traffic scene understanding. In: BMVC (2009)
5. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: the KITTI dataset. Int. J. Robot. Res. **32**, 1231–1237 (2013)
6. Cordts, M., et al.: The cityscapes dataset for semantic urban scene understanding. In: CVPR (2016)
7. Menze, M., Geiger, A.: Object scene flow for autonomous vehicles. In: CVPR (2015)
8. Franke, U., et al.: Making bertha see. In: ICCV Workshops. IEEE (2013)
9. Xiang, Y., Fox, D.: DA-RNN: semantic mapping with data associated recurrent neural networks. In: Robotics: Science and Systems (RSS) (2017)

10. Kundu, A., Li, Y., Dellaert, F., Li, F., Rehg, J.M.: Joint semantic segmentation and 3D reconstruction from monocular video. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8694, pp. 703–718. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10599-4_45
11. Szegedy, C., et al.: Going deeper with convolutions. In: CVPR (2015)
12. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: CVPR (2016)
13. Szegedy, C., Ioffe, S., Vanhoucke, V.: Inception-v4, inception-resnet and the impact of residual connections on learning. CoRR (2016)
14. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: CVPR (2017)
15. Chollet, F.: Xception: deep learning with depthwise separable convolutions. In: CVPR (2017)
16. Howard, A.G., et al.: MobileNets: efficient convolutional neural networks for mobile vision applications. arXiv preprint [arXiv:1704.04861](https://arxiv.org/abs/1704.04861) (2017)
17. Zhang, X., Zhou, X., Lin, M., Sun, J.: ShuffleNet: an extremely efficient convolutional neural network for mobile devices. In: CVPR (2018)
18. Yu, F., Koltun, V.: Multi-scale context aggregation by dilated convolutions. In: ICLR (2016)
19. Yu, F., Koltun, V., Funkhouser, T.: Dilated residual networks. In: CVPR (2017)
20. Paszke, A., Chaurasia, A., Kim, S., Culurciello, E.: ENet: a deep neural network architecture for real-time semantic segmentation. arXiv preprint [arXiv:1606.02147](https://arxiv.org/abs/1606.02147) (2016)
21. Romera, E., Alvarez, J.M., Bergasa, L.M., Arroyo, R.: ERFNet: efficient residual factorized convnet for real-time semantic segmentation. *IEEE Trans. Intell. Transp. Syst.* **19**, 263–272 (2018)
22. Jin, J., Dundar, A., Culurciello, E.: Flattened convolutional neural networks for feedforward acceleration. arXiv preprint [arXiv:1412.5474](https://arxiv.org/abs/1412.5474) (2014)
23. Chen, W., Wilson, J., Tyree, S., Weinberger, K., Chen, Y.: Compressing neural networks with the hashing trick. In: ICML (2015)
24. Han, S., Mao, H., Dally, W.J.: Deep compression: compressing deep neural networks with pruning, trained quantization and Huffman coding. In: ICLR (2016)
25. Wu, J., Leng, C., Wang, Y., Hu, Q., Cheng, J.: Quantized convolutional neural networks for mobile devices. In: CVPR (2016)
26. Zhao, H., Qi, X., Shen, X., Shi, J., Jia, J.: ICNet for real-time semantic segmentation on high-resolution images. arXiv preprint [arXiv:1704.08545](https://arxiv.org/abs/1704.08545) (2017)
27. Jaderberg, M., Vedaldi, A., Zisserman, A.: Speeding up convolutional neural networks with low rank expansions. In: BMVC (2014)
28. Rastegari, M., Ordonez, V., Redmon, J., Farhadi, A.: XNOR-Net: imagenet classification using binary convolutional neural networks. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9908, pp. 525–542. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46493-0_32
29. Hwang, K., Sung, W.: Fixed-point feedforward deep neural network design using weights 1, 0, and -1 . In: 2014 IEEE Workshop on Signal Processing Systems (SiPS) (2014)
30. Courbariaux, M., Hubara, I., Soudry, D., El-Yaniv, R., Bengio, Y.: Binarized neural networks: training neural networks with weights and activations constrained to $+1$ or -1 . arXiv preprint [arXiv:1602.02830](https://arxiv.org/abs/1602.02830) (2016)
31. Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., Bengio, Y.: Quantized neural networks: training neural networks with low precision weights and activations. arXiv preprint [arXiv:1609.07061](https://arxiv.org/abs/1609.07061) (2016)

32. Liu, B., Wang, M., Foroosh, H., Tappen, M., Pensky, M.: Sparse convolutional neural networks. In: CVPR, pp. 806–814 (2015)
33. Wen, W., Wu, C., Wang, Y., Chen, Y., Li, H.: Learning structured sparsity in deep neural networks. In: NIPS, pp. 2074–2082 (2016)
34. Bagherinezhad, H., Rastegari, M., Farhadi, A.: LCNN: lookup-based convolutional neural network. In: CVPR (2017)
35. Holschneider, M., Kronland-Martinet, R., Morlet, J., Tchamitchian, P.: A real-time algorithm for signal analysis with the help of the wavelet transform. In: Combes, J.M., Grossmann, A., Tchamitchian, P. (eds.) *Wavelets*, pp. 286–297. Springer, Heidelberg (1990). https://doi.org/10.1007/978-3-642-75988-8_28
36. Mehta, S., Mercan, E., Bartlett, J., Weaver, D.L., Elmore, J.G., Shapiro, L.G.: Learning to segment breast biopsy whole slide images. In: WACV (2018)
37. Wang, P., et al.: Understanding convolution for semantic segmentation. In: WACV (2018)
38. Graves, A., Fernández, S., Schmidhuber, J.: Multi-dimensional recurrent neural networks. In: de Sá, J.M., Alexandre, L.A., Duch, W., Mandic, D. (eds.) *ICANN 2007*. LNCS, vol. 4668, pp. 549–558. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74690-4_56
39. Badrinarayanan, V., Kendall, A., Cipolla, R.: SegNet: a deep convolutional encoder-decoder architecture for image segmentation. *TPAMI* **39**, 2481–2495 (2017)
40. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) *MICCAI 2015*. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
41. Hariharan, B., Arbeláez, P., Girshick, R., Malik, J.: Hypercolumns for object segmentation and fine-grained localization. In: CVPR (2015)
42. Dai, J., He, K., Sun, J.: Convolutional feature masking for joint object and stuff segmentation. In: CVPR (2015)
43. Caesar, H., Uijlings, J., Ferrari, V.: Region-based semantic segmentation with end-to-end training. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016*. LNCS, vol. 9905, pp. 381–397. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46448-0_23
44. Lin, G., Milan, A., Shen, C., Reid, I.: RefineNet: multi-path refinement networks for high-resolution semantic segmentation. In: CVPR (2017)
45. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR (2015)
46. Noh, H., Hong, S., Han, B.: Learning deconvolution network for semantic segmentation. In: ICCV (2015)
47. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
48. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks. In: NIPS (2012)
49. Ioffe, S., Szegedy, C.: Batch normalization: accelerating deep network training by reducing internal covariate shift. In: ICML (2015)
50. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: surpassing human-level performance on imagenet classification. In: ICCV (2015)
51. Neuhold, G., Ollmann, T., Rota Bulò, S., Kotschieder, P.: The mapillary vistas dataset for semantic understanding of street scenes. In: ICCV (2017)
52. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (VOC) challenge. *IJCV* **88**, 303–338 (2010)

53. Hariharan, B., Arbeláez, P., Bourdev, L., Maji, S., Malik, J.: Semantic contours from inverse detectors. In: ICCV (2011)
54. Lin, T.-Y., et al.: Microsoft COCO: common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8693, pp. 740–755. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10602-1_48
55. Iandola, F.N., Han, S., Moskewicz, M.W., Ashraf, K., Dally, W.J., Keutzer, K.: SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <0.5 MB model size. arXiv preprint [arXiv:1602.07360](https://arxiv.org/abs/1602.07360) (2016)
56. Yasin, A., Ben-Asher, Y., Mendelson, A.: Deep-dive analysis of the data analytics workload in cloudsuite. In: 2014 IEEE International Symposium on Workload Characterization (IISWC) (2014)
57. Wu, Y., Wang, Y., Pan, Y., Yang, C., Owens, J.D.: Performance characterization of high-level programming models for GPU graph analytics. In: 2015 IEEE International Symposium on Workload Characterization (IISWC), pp. 66–75. IEEE (2015)
58. PyTorch: Tensors and dynamic neural networks in python with strong GPU acceleration. <http://pytorch.org/>. Accessed 08 Feb 2018
59. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. In: ICLR (2015)
60. NVPROF: CUDA toolkit documentation. <http://docs.nvidia.com/cuda/profiler-users-guide/index.html>. Accessed 08 Feb 2018
61. TegraTools: NVIDIA embedded computing. <https://developer.nvidia.com/embedded/develop/tools>. Accessed 08 Feb 2018
62. PowerTop: For PowerTOP saving power on IA isn't everything. It is the only thing! <https://01.org/powertop/>. Accessed 08 Feb 2018
63. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: ICLR (2015)
64. Trembl, M., et al.: Speeding up semantic segmentation for autonomous driving. In: MLITS, NIPS Workshop (2016)
65. Ghiasi, G., Fowlkes, C.C.: Laplacian pyramid reconstruction and refinement for semantic segmentation. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9907, pp. 519–534. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46487-9_32
66. Mehta, S., Mercan, E., Bartlett, J., Weaver, D., Elmore, J., Shapiro, L.: Y-Net: joint segmentation and classification for diagnosis of breast biopsy images. In: MICCAI (2018)
67. Nair, V., Hinton, G.E.: Rectified linear units improve restricted Boltzmann machines. In: ICML (2010)