

STACKED U-NETS: A NO-FRILLS APPROACH TO NATURAL IMAGE SEGMENTATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Many imaging tasks require global information about all pixels in an image. Conventional bottom-up classification networks globalize information by decreasing resolution; features are pooled and down-sampled into a single output. But for semantic segmentation and object detection tasks, a network must provide higher-resolution pixel-level outputs. To globalize information while preserving resolution, many researchers propose the inclusion of sophisticated auxiliary blocks, but these come at the cost of a considerable increase in network size and computational cost. This paper proposes stacked u-nets (SUNets), which iteratively combine features from different resolution scales while maintaining resolution. SUNets leverage the information globalization power of u-nets in a deeper network architectures that is capable of handling the complexity of natural images. SUNets perform extremely well on semantic segmentation tasks using a small number of parameters.

1 INTRODUCTION

Semantic segmentation methods decompose an image into groups of pixels, each representing a common object class. While the output of a segmentation contains object labels assigned at the local (pixel) level, each label depends on global information about the image, such as textures, colors, and object boundaries that may span large regions. Simple image classification algorithms consolidate global information by successively pooling features until the final output is a single label containing information from the entire image. In contrast, segmentation methods must output a full-resolution map of labels. Thus, a successful segmentation method must address this key question: how can we learn long-distance contextual information while at the same time retaining high spatial resolution at the output for identifying small objects and sharp boundaries?

For natural image processing, most research has answered this question using one of two approaches. One approach is to use very few pooling layers, thus maintaining resolution (although methods may still require a small number of deconvolution layers (Lin et al., 2017a; Fu et al., 2017; Jégou et al., 2017; Ghiasi & Fowlkes, 2016)). Large fields of view are achieved using dilated convolutions, which span large regions. By maintaining resolution at each layer, this approach preserves substantial amounts of signal about smaller and less salient objects. However, this is achieved at the cost of computationally expensive and memory exhaustive training/inference. The second related approach is to produce auxiliary context aggregation blocks (Krähenbühl & Koltun, 2011; Chen et al., 2016a; Chandra & Kokkinos, 2016; Chandra et al., 2017; Zheng et al., 2015; Yu & Koltun, 2015; Zhao et al., 2017) that contain features at different distance scales, and then merge these blocks to produce a final segmentation. This category includes many well-known techniques such as dense CRF (Krähenbühl & Koltun, 2011) (conditional random fields) and spatial pyramid pooling (Chen et al., 2016a).

These approaches suffer from the following challenges:

1. Deconvolutional (i.e., encoder-decoder) architectures perform significant nonlinear computation at low resolutions, but do very little processing of high-resolution features. During the convolution/encoding stage, pooling layers can move information over large distances, but information about small objects is often lost. During the deconvolution/decoding stage, high- and low-resolution features are merged to produce upsampled feature maps. However, the high-resolution

inputs to deconvolutional layers come from relatively shallow layers that do not effectively encode semantic information.

2. Image classification networks are parameter heavy (44.5M parameters for ResNet-101), and segmentation methods built on top of these classification networks are often even more burdensome. For example, on top of the resnet-101 architecture, PSPNet (Zhao et al., 2017) uses 22M *additional* parameters for context aggregation, while the ASPP and Cascade versions of the Deeplab network utilize 14.5M (Chen et al., 2016a) and 40M (Chen et al., 2017) *additional* parameters, respectively.

A popular and simple approach to segmentation is u-nets, which perform a chain of convolutional/downsampling operations, followed by a chain of deconvolutional/upsampling layers that see information from both low- and high-resolution scales. These u-net architectures are state-of-the-art for medical image segmentation (Ronneberger et al., 2015), but they do not perform well when confronted with the complex color profiles, lighting effects, and occlusions present in natural images.

We expand the power of u-nets by stacking u-net blocks into deep architectures. This addresses the two challenges discussed above: As data passes through multiple u-net blocks, high-resolution features are mixed with low-resolution context information and processed through many layers to produce informative high-resolution features. Furthermore, stacked U-net models require fewer feature maps per layer than conventional architectures, and thus achieve higher performance with far fewer parameters. Our smallest model exceeds the performance of ResNet-101 on the PASCAL VOC 2012 semantic segmentation task by 4.5% mIoU, while having $\sim 7\times$ fewer parameters.

2 RELATED WORK

Many models (Chen et al., 2017; Zhao et al., 2017; Peng et al., 2017; Chen et al., 2016a; Wang et al., 2017; Ghiasi & Fowlkes, 2016; Lin et al., 2017a; Fu et al., 2017) have boosted the performance of semantic segmentation networks. These gains are mainly attributed to the use of pre-trained models, dilated convolutional layers (Chen et al., 2016a; Yu & Koltun, 2015) and fully convolutional architectures (DCNN) (Long et al., 2015). These works employ a range of strategies to tap contextual information, which fall into three major categories.

Context Aggregation Modules: These architectures place a special module on top of a pre-trained network that integrates context information at different distance scales. The development of fast and efficient algorithm for DenseCRF (Krähenbühl & Koltun, 2011) led to the development of numerous algorithms (Chen et al., 2016a; Liang-Chieh et al., 2015; Chandra & Kokkinos, 2016; Chandra et al., 2017) incorporating it on top of the output belief map. Moreover, the joint training of CRF and CNN parameters was made possible by Zheng et al. (2015); Schwing & Urtasun (2015). ParseNet (Liu et al., 2015) exploits image-level feature information at each layer to learn global contextual information. In contrast, Zhao et al. (2017); Chen et al. (2017; 2016a) realized substantial performance improvements by employing parallel layers of spatial pyramid pooling.

Image Pyramid: The networks proposed in Lin et al. (2016); Chen et al. (2016b) learn context information by simultaneously processing inputs at different scales and merging the output from all scales. Recently, Dai et al. (2017) developed a deformable network that adaptively determines an object’s scale and accordingly adjusts the receptive field size of each activation function. On the other hand, Singh & Davis (2017) proposed a new training paradigm for object detection networks that trains each object instance only using the proposals closest to the ground truth scale.

Encoder-Decoder: These models consist of an encoder network and one or many blocks of decoder layers. The decoder fine-tunes the pixel-level labels by merging the contextual information from feature maps learned at all the intermediate layers. Usually, a popular bottom-up pre-trained classification network such as ResNet (He et al., 2016a), VGG (Simonyan & Zisserman, 2014) or DenseNet (Huang et al., 2017) serves as an encoder model. U-net (Ronneberger et al., 2015) popularly employs skip connections between an encoder and its corresponding decoding layers. Ghiasi & Fowlkes (2016) use a Laplacian pyramid reconstruction network to selectively refine the low resolution maps. Refinenet (Lin et al., 2017a) employs sophisticated decoder modules at each scale on top of the ResNet encoder, while Chen et al. (2018) utilize a simple two-level decoding of feature maps from the Xception network (Chollet, 2017). In short, the structure in Chen et al. (2018); Ghiasi & Fowlkes (2016) is a hybrid of decoding and context aggregation modules.

2.1 USE OF PRE-TRAINED NETS

Many of the networks described above make extensive use of image classification networks that were pre-trained for other purposes. Although parameter heavy, much fundamental work on segmentation (FCN (Long et al., 2015), dilated nets (Yu & Koltun, 2015), u-nets (Ronneberger et al., 2015) and CRF (Zheng et al., 2015)) was built on VGG. These architectures share common origins in that they were designed for the ImageNet competition and features are processed bottom-up. This prototype works well when the network has to identify only a single object without exact pixel localization. However, when extended to localization tasks such as segmentation and object detection, it is not clear whether the complete potential of these networks has been properly tapped. Recent work on object detection (Singh & Davis, 2017) also echoes a similar concern.

3 U-NETS REVISITED

The original u-net architecture was introduced by Ronneberger et al. (2015), and produced almost perfect segmentation of cells in biomedical images using very little training data. The structure of u-nets makes it possible to capture context information at multiple scales and propagate them to the higher resolution layers. These higher order features have enabled u-nets to outperform previous deep models on various tasks including semantic segmentation (Milletari et al., 2016), depth-fusion (Riegler et al., 2017), image translation (Isola et al., 2017) and human-pose estimation (Newell et al., 2016). Moreover, driven by the initial success of u-nets, many recent works on semantic segmentation (Lin et al., 2017a; Ghiasi & Fowlkes, 2016; Fu et al., 2017) and object detection (Shrivastava et al., 2016; Lin et al., 2017b) also propose an encoder-decoder deep architecture.

The u-net architecture evenly distributes its capacity among the encoder and decoder modules. Moreover, the complete network can be trained in an end-to-end setting. In contrast, the more recent architectures reviewed in Section 2 do not equally distribute the processing of top-down and bottom-up features. Since these architectures are built on top of pre-trained feature extractors (Simonyan & Zisserman, 2014; He et al., 2016a), the decoder modules are trained separately and sometimes in multiple stages. To overcome these drawbacks, Jégou et al. (2017) proposed an equivalent u-net based on the Densenet (Huang et al., 2017) architecture. However, DenseNet is memory intensive, and adding additional decoder layers leads to a further increase in memory usage. Given these drawbacks, the effectiveness of these architectures on different datasets and applications is unclear.

The goal of this paper is to realize the benefits of u-nets (small size, easy trainability, high performance) for complex natural image segmentation problems. Specifically, we propose a new architecture composed of multiple stacks of u-nets. The network executes repeated processing of both top-down as well as bottom-up features and captures long-distance spatial information at multiple resolutions. The network is trained end-to-end on image classification tasks and can be seamlessly applied to semantic segmentation without any auxiliary modules on top.

Our stacked u-net (SUNet) architecture shares some similarity with other related stacked encoder-decoder structures. Fu et al. (2017) use multiple stacks (upto 3) of de-convolutional networks on top of a powerful encoder (DenseNet) while Newell et al. (2016) apply multiple stacks of u-net modules for human-pose estimation. However, the processing of features inside each u-net module in Newell et al. (2016) differs from ours. Newell et al. (2016) replace each convolutional block with a residual module, utilizes nearest-neighbor upsampling for deconvolution and operate at fix resolution. In contrast, SUNets retain the basic u-net structure from Ronneberger et al. (2015), operates without any intermediate supervision and processes features by progressively downsampling.

3.1 U-NET MODULE IMPLEMENTATION

Figure 1 illustrates the design of the u-net module employed in our stacked architecture. Each module is composed of 10 pre-activated convolutional blocks each preceded by a batch-normalization and a ReLU non-linearity. The pooling/unpooling operation, handled by the strided convolutional/deconvolutional layers, facilitates information exchange between the lower and the higher resolution features. A skip connection branches off at the output of the first encoder block, $E1$. Following this, the $E2$ and $D2$ blocks capture long distance context information using lower resolution feature maps and merge the information back with the high resolution features from $E1$ at the output of $D2$. Every layer (except for the bottleneck layers) uses 3×3 kernels, and outputs a




Figure 1: A typical u-net module with outer residual connection. M is the number of input features. Across the u-net module, each layer has the same number of output feature maps (except for the final 1×1 filter), which we denote N . For better understanding the figure also includes the field of view (FoV) of each convolutional kernel (top) and the feature map size at the output of each filter (bottom), assuming a 64×64 input I . Best viewed in color.

fixed number of feature maps, N . To mitigate high frequency noise from the sampling operation, each strided conv/de-conv layer is followed by a convolution. Unlike traditional u-nets, the design of convolutional layers in our u-net module helps in retaining the original size of the feature maps at its output. Consequently, multiple u-net modules can be stacked without losing resolution.

In the following, we briefly highlight some of the design choices of the architecture. In comparison to traditional u-nets, the max-pooling operation is replaced with strided convolution for SUNets. The use of strided convolutions enables different filters in each u-net module to operate at different resolutions (see the discussion in Section 5). Moreover, the repeated use of max-pooling operations can cause gridding artifacts in dilated networks (Yu et al., 2017).

Unlike the u-nets of Ronneberger et al. (2015); Newell et al. (2016), our u-net module is comprised of only two levels of depth. We considered two major factors in choosing the depth: field of view (FoV) of the innermost conv filter and the total number of parameters in a single u-net module. The number of parameters influences the total number of stacks in SUNets. While keeping the total parameters of SUNet approximately constant, we experimented with a higher depth of three and four. We found that the increase in depth indeed led to a decline in performance for the image classification task. This may not be surprising, given that a SUNet with depth of two is able to stack more u-net modules. Moreover, deeper u-net modules make it harder to train the inner-most convolutional layers due to the vanishing gradients problem (Bengio et al., 1994). For instance, in our current design, the maximum length of the gradient path is six. The popular classification networks of He et al. (2016a); Huang et al. (2017) are known to operate primarily on features with 28^2 and 14^2 resolution. At this scale, the effective FoV of 19 is more than sufficient to capture long-distance contextual information. Moreover, the stacking of multiple u-net modules will also serve to increase the effective FoV of higher layers.

SUNets train best when there is sufficient gradient flow to the bottom-most u-net layers. To avoid vanishing gradients, we include a skip connection (He et al., 2016a; Huang et al., 2017) around each u-net module. Also, inspired by the design of bottleneck blocks (He et al., 2016a), we also include 1×1 convolutional layers. Bottleneck layers restricts the number of input features to a small number (N), avoiding parameter inflation.

When stacking multiple u-nets it makes sense for each u-net module to reuse the raw feature maps from all the preceding u-net modules. Thus we also explored replacing the identity connection with dense connectivity (Huang et al., 2017). This new network is memory intensive which in turn prevented proper learning of the batch-norm parameters. Instead, we chose to utilize dense connectivity only within each u-net, i.e., while reusing feature maps from $E1$ at $D1$. Thus the proposed u-net module leverages skip connectivity without getting burdened.

4 SUNETS: STACKED U-NETS FOR CLASSIFICATION

Layers	Output Size	SUNet-64	SUNet-128	SUNet-7-128
Convolution	112×112	7×7 conv, 64, stride 2		
Residual Block	56×56	$\begin{bmatrix} 3 \times 3 \text{ conv, 128, stride 2} \\ 3 \times 3 \text{ conv, 128, stride 1} \end{bmatrix} \times 1$		
UNets Block (1)	56×56	$\begin{bmatrix} 1 \times 1 \text{ conv, 64} \\ \text{U-Net, N=64} \\ 1 \times 1 \text{ conv, 256} \end{bmatrix} \times 2$	$\begin{bmatrix} 1 \times 1 \text{ conv, 128} \\ \text{U-Net, N=128} \\ 1 \times 1 \text{ conv, 512} \end{bmatrix} \times 2$	$\begin{bmatrix} 1 \times 1 \text{ conv, 128} \\ \text{U-Net, N=128} \\ 1 \times 1 \text{ conv, 512} \end{bmatrix} \times 2$
Transition Layer	28×28	2×2 average pool, stride 2		
UNets Block (2)	28×28	$\begin{bmatrix} 1 \times 1 \text{ conv, 64} \\ \text{U-Net, N=64} \\ 1 \times 1 \text{ conv, 512} \end{bmatrix} \times 4$	$\begin{bmatrix} 1 \times 1 \text{ conv, 128} \\ \text{U-Net, N=128} \\ 1 \times 1 \text{ conv, 1024} \end{bmatrix} \times 4$	$\begin{bmatrix} 1 \times 1 \text{ conv, 128} \\ \text{U-Net, N=128} \\ 1 \times 1 \text{ conv, 1280} \end{bmatrix} \times 7$
Transition Layer	14×14	2×2 average pool, stride 2		
UNets Block (3)	14×14	$\begin{bmatrix} 1 \times 1 \text{ conv, 64} \\ \text{U-Net, N=64} \\ 1 \times 1 \text{ conv, 768} \end{bmatrix} \times 4$	$\begin{bmatrix} 1 \times 1 \text{ conv, 128} \\ \text{U-Net, N=128} \\ 1 \times 1 \text{ conv, 1536} \end{bmatrix} \times 4$	$\begin{bmatrix} 1 \times 1 \text{ conv, 128} \\ \text{U-Net, N=128} \\ 1 \times 1 \text{ conv, 2048} \end{bmatrix} \times 7$
Transition Layer	7×7	2×2 average pool, stride 2		
UNets Block (4)	7×7	$\begin{bmatrix} 1 \times 1 \text{ conv, 64} \\ \text{U-Net}^+, \text{N=64} \\ 1 \times 1 \text{ conv, 1024} \end{bmatrix} \times 1$	$\begin{bmatrix} 1 \times 1 \text{ conv, 128} \\ \text{U-Net}^+, \text{N=128} \\ 1 \times 1 \text{ conv, 2048} \end{bmatrix} \times 1$	$\begin{bmatrix} 1 \times 1 \text{ conv, 128} \\ \text{U-Net}^+, \text{N=128} \\ 1 \times 1 \text{ conv, 2304} \end{bmatrix} \times 1$
Classification Layer	1×1	7×7 global average pool 1000D fully-connected, softmax		
Total Layers		110	110	170
Params		6.9M	24.6M	37.7M

Table 1: SUNet architectures for ImageNet. N denotes the number of filters per convolutional layer. Note that the building block in bracket refers to the integrated u-net module shown in Figure 1.

Before addressing segmentation, we describe a stacked u-net (SUNet) architecture that is appropriate for image classification. Because the amount of labeled data available for classification is much larger than for segmentation, classification tasks are often used to pre-train feature extraction networks, which are then adapted to perform segmentation.

The network design of SUNets for ImageNet classification is summarized in Table 1. Note that each ‘‘conv’’ layer shown in the table corresponds to a sequence of ‘‘BN-ReLU-Conv’’ layers. The three listed configurations mainly differ in the number of output feature maps N of each convolutional layer and the total number of stacks in blocks 2 and 3. Input images are processed using a 7×7 conv filter followed by a residual block. Inspired by the work on dilated resnet (Yu et al., 2017), the conventional max-pooling layer at this stage is replaced by a strided convolutional layer inside the residual block. Subsequently, the feature maps are processed bottom-up as well as top-down by multiple stacks of u-nets at different scales and with regularly decreasing resolutions. The feature map input size to block 4 is 7×7 and is further reduced to 2×2 at the input to the encoder $E2$ of the u-net module. At this resolution, it is not possible to have $E2$ and $D2$ layers, and hence a trimmed version of u-nets (u-net⁺) are employed in block 4. The u-net⁺ includes a single level of encoder and decoder ($E1, D1$) processing. Following block 4, a global average pooling is performed on features and transferred to a softmax classifier.

The residual connection in all but the first u-net in each block is implemented as an identity mapping. In the first u-nets the skip connection is implemented using an expansion layer i.e., a 1×1 conv filter. The number of feature map outputs from each block approximately equates to the total number of feature maps generated by all the preceding u-net modules. This arrangement allows flexibility for the network to retain all the raw feature maps of the preceding modules. Moreover among all other possibilities, the above architectures were picked because their performance on the image classification task is roughly equivalent to the ResNet-18, 50 and 101 network architectures (discussed in Section 6), albeit with fewer parameters.




Figure 2: Dilated SUNet-7-128 network for segmentation at $output_stride = 8$. For dilation > 1 , the feature maps are processed with a varying range of dilation factors inside each u-net module (see inset). The de-gridding filters smooth out aliasing artifacts that occur during deconvolution.

As in ResNet and DenseNet, most of the processing in SUNet is performed at the feature scale of 14×14 (46 conv layers) and 7×7 (44 conv layers). However, the order at which the local information is processed can lead to a substantial gap in performance between ResNet and SUNet when extended to object localization, detection, and image segmentation tasks. All these task demands pixel-level localization and hence require a deep architecture that can efficiently integrate local and global cues. The development of SUNet is a first step towards achieving this objective. Intuitively, multiple stacks of u-nets can be seen as multiple iterations of the message passing operation in a CRF.

5 DILATED SUNETS FOR SEGMENTATION

One can directly extend SUNet to segmentation by simply removing a global average pooling layer and operating the network in fully convolutional mode. Akin to other works on semantic segmentation (Chen et al., 2017), the output feature maps are rescaled using bilinear interpolation to the input image size before passing into the softmax layer with multi-class cross-entropy loss.

5.1 DILATION

For an input image of 512×512 , the output map size at the softmax is 16×16 i.e., subsampled by a factor of 32. This is insufficient to preserve precise pixel-level localization information at its output. The precision can be improved by increasing the output map size of the network. This is realized by dropping the pooling stride at the transition layer. Merely eliminating stride leads to the reduction in the receptive field of the subsequent layers by a factor of two. Consequently, this reduces the influence of long-distance context information on the output prediction. Nevertheless, the receptive field is restored to that of the original network by operating each convolutional filter in the subsequent layers at a dilation factor of 2 (Chen et al., 2016a; Yu & Koltun, 2015).

5.2 MULTIGRID

Figure 2 shows a sample dilated SUNet architecture used for the semantic segmentation task. Similar to (Chen et al., 2017), we define $output_stride$ to be the ratio of resolution of an input image to that of its output feature map.

To sample at an $output_stride$ of 8 the pooling layers preceding blocks (3) and (4) are discarded. Following this, the dilation factor for each u-net module in blocks 3 and 4 is fixed at 2 and 4, respectively. In each subsequent u-net module the 3×3 conv layers are operated with $stride = 1$. To keep the receptive field of the low-resolution layers in these modules constant, a dilation is applied. This arrangement facilitates the network to preserve spatial information learned from the prior modules (because there is no downsampling in the final u-net block) while preserving the distance scale of features within each block. As an example, the inset in Figure 2 displays the effective dilation rate for each layer in the u-net module at block 3. Similarly, the dilation rate of each layer (except for bottleneck layers) in the u-net⁺ module will be twice that of the corresponding layers in block 3. The steady increase and decrease of dilation factors inside each u-net module is analogous to multigrid solvers for linear systems (Brandt, 1977; Briggs et al., 2000), which use grids

at different scales to move information globally. Many recent works (Dai et al., 2017; Wang et al., 2017; Chen et al., 2017) on deep networks advocate the use of special structures for information globalization. In SUNet, the multigrid structure is baked into the model, and no further “frills” are needed to globalize information.

5.3 DE-GRIDDING FILTERS

By adopting dilated SUNets, we observe a vast improvement in segmentation performance. However, for $output_stride = 8$ the segmentation map displays gridding artifacts (Wang et al., 2017; Yu et al., 2017). This aliasing artifact is introduced when the sampling rate of the dilated layer is lower than the high-frequency content of input feature maps. The final conv filter of u-net⁺ operates at the dilation factor of 4. Directly transferring u-net⁺’s feature map output to a classification layer can cause gridding artifacts. Following Yu et al. (2017), the u-net⁺ module is followed by two layers of de-gridding filters with progressively decreasing dilation factor. Each filter is a 3×3 conv layer and outputs 512 feature maps.

In all, SUNet does not require any additional post-hoc structural changes popularized by recent works such as decoding layers (Lin et al., 2017a), appending context aggregation blocks (Chen et al., 2017; Zhao et al., 2017) and learning conditional random fields (Chandra et al., 2017). Hence we regard SUNet as a “no-frills” network.

6 EXPERIMENTS

6.1 IMAGENET CLASSIFICATION

In this section, we evaluate three SUNet architectures on the ILSVRC-2012 classification dataset, which contains 1.28M training images and 50,000 images for validation, with labels distributed over 1000 classes. Training utilized the same data augmentation scheme used for ResNet and DenseNet. Following common practice (He et al., 2016a;b), we apply a 224×224 center crop on test images and report Top-1 and Top-5 error on the validation set. More details in Appendix A.

Table 2 compares the performance of SUNet against other classification networks. The comparison is restricted to only ResNet and DenseNet models as most recent work on segmentation builds on top of them. The notable point about the result is that the repeated top-down and the bottom-up processing of features performs equivalently to state-of-the-art classification networks.

We emphasize here that our objective is not to surpass classification accuracy but instead to build a better architecture for segmentation by pre-training on a classification task. Indeed, each SUNet model was selected such that it is the counterpart for the corresponding ResNet model.

Model	Top-1	Top-5	Depth	Params
ResNet-18 [†]	30.24	10.92	18	11.7M
ResNet-50 [†]	23.85	7.13	50	25.6M
ResNet-101 [†]	22.63	6.44	101	44.5M
DenseNet-201 [†]	22.80	6.43	201	20M
DenseNet-161 [†]	22.35	6.20	161	28.5M
SUNet-64	29.28	10.21	111	6.9M
SUNet-128	23.64	7.42	111	24.6M
SUNet-7-128	22.47	6.85	171	37.7M

Table 2: Error rates for classification networks on the ImageNet 2012 validation set. ‘†’ denotes error rates from the official PyTorch implementation.

6.2 SEMANTIC SEGMENTATION

Semantic segmentation networks were built using the dilated version of the ImageNet pre-trained SUNet models (Section 5). We evaluate on the PASCAL VOC 2012 semantic segmentation benchmark (Everingham et al., 2015) and urban scene understanding Cityscape (Cordts et al., 2016)

Model	mIoU
ResNet-101 (Chen et al., 2017)	68.39
SUNet-64	72.85
SUNet-128	77.16
SUNet-7-128	78.95

Table 3: The semantic segmentation performance of dilated SUNet and ResNet-101 networks on PASCAL VOC 2012 validation set trained with $output_stride = 16$. Relative to the ResNet-101 network, all SUNets perform very well.

datasets. The performance on each of these datasets is reported using intersection-over-union (IoU) averaged over all classes. The experimental details are shared in Appendix B.

Unless mentioned, for all our experiments we set $output_stride = 16$. This means only the u-net modules in the final block (4) operate at dilation factor of two; all other modules use the same stride as in the original classification model. The training and inference are $2\times$ faster with $output_stride = 16$ rather than 8.

6.2.1 ABLATION STUDY

We experiment with different SUNet variants on the PASCAL VOC 2012 dataset.

SUNets vs ResNet-101: We compare the performance of the dilated SUNet architecture on semantic segmentation against the popular dilated ResNet-101 model. Models were fine-tuned on the “trainaug” set without the degridding layers and evaluated on the validation set.

The performance of the plain dilated SUNets surpasses that of ResNet-101 by a wide margin (Table 3). In fact, the smallest SUNet model, SUNet-64 with 6.7M parameters, beats ResNet-101 (with 44.5M) by an absolute margin of 4.5% IoU while SUNet-7-128, the counterpart network to ResNet-101, improves by over 10.5% IoU. This is substantial, given that the gap between the ResNet-101 and VGG-16 models is $\sim 3\%$ (Chen et al., 2016a) (at $output_stride = 8$). This contrasts with the negligible performance differences observed on classification, and suggests that specialized segmentation network architectures can surpass architectures adapted from classification.

Finally, we note that, although SUNets were designed for pixel-level localization tasks, the selected models were chosen only based on their classification performance. By linking the model selection process to the primary task (segmentation) there is a possibility of improving performance.

Multigrid vs Downsampling: We compare the performance of multigrid dilation (as shown in Figure 2) inside each u-net against the usual downsampling (Figure 1). We consider a dilated SUNet-7-128 network and report performance at different training $output_stride$. The result is summarized in Table 4. For a dilated network, replacing strided convolutional layers with the corresponding dilated layers is more logical as well as beneficial. This is because, when operating dilated convolutional layers with $stride > 1$, alternate feature points are dropped without being processed by any of the filters, leading to high frequency noise at the decoder output. Furthermore, due to a skip connection, the features from the lower layers are also corrupted. Due to error propagation, this effect is more prominent in a network with many dilated modules (for eg., $output_stride = 8$).

Output Stride and Inference Strategy: Finally, we experiment with three different training output strides (8,16,32) and multi-scale inference at test time. For OS = 32, none of the layers are dilated and hence de-gridding layers were not used. Training with an OS = 32 is equivalent to fine-tuning a classification network with a global pooling layer removed. For multi-scale inference, each input image is scaled and tested using multiple scales $\{0.5, 0.75, 1.0, 1.25\}$ and its left-right flipped image. The average over all output maps is used in the final prediction. See results in Table 5. We note that:

1. The network trained with OS = 32 performs 0.7 IoU better (with single scale) than the Resnet-101 and Resnet-152 models (Wu et al., 2016) each trained at OS = 8. This is significant as SUNet model outputs $16\times$ fewer pixels at $4\times$ faster training/inference with no performance drop.
2. The degridding layers do not improve performance at OS = 16. Since there is only small change in dilation factor between the final SUNet layer & classification layer, aliasing is not problematic.
3. The margin of performance improvement decreases with decrease in training OS. Given this and the above fact, subsequently we only report performance for models trained at OS = 16 without any degridding layers.

Pretraining with MS-COCO: Following common practice (Chen et al., 2017), we pretrain SUNet-7-128 with the MS-COCO dataset (Lin et al., 2014). The dataset contains pixel-level annotation for 80 object classes. Except for the PASCAL VOC classes, the pixel annotation for all other classes is set to the background class. Following this, the images having < 1000 annotated pixel labels were dropped yielding 90,000 training images. After pretraining, the model is fine-tuned on “trainaug” for 5K iterations with $10\times$ smaller initial learning rate. In the end, our model achieves 83.27% mIoU on the validation set. This performance is slightly better than the ResNet+ASPP model (Chen et al., 2017) (82.70%) and equivalent to Xception+ASPP+Decoder model of Chen et al. (2018) (83.34%).

OS	Strided conv	Multigrid
8	65.99	78.64
16	78.25	78.95

Table 4: Performance comparison of multigrid dilation against strided convolution inside each u-net module, using the SUNet-7-128 model and evaluated using mean IoU. **OS** denotes *output_stride* during training.

train OS	eval OS	DL	MS	Flip	mIoU
32	32				76.03
32	32		✓		77.58
32	32		✓	✓	77.57
16	16				78.95
16	16	✓			78.10
16	16		✓		80.22
16	16		✓	✓	80.40
8	8				78.64
8	8	✓			78.88
8	8		✓		80.37
8	8		✓	✓	80.50

Table 5: Performance comparison at various *output_stride* and inference strategies. **MS**: Multi-scale, **DL**: with Degridding Layers

Methods	mIoU
Piecewise (VGG16) (Lin et al., 2016)	78.0
LRR+CRF (Ghiasi & Fowlkes, 2016)	77.3
DeepLabv2+CRF (Chen et al., 2016a)	79.7
Large-Kernel+CRF (Peng et al., 2017)	82.2
Deep Layer Cascade* (Li et al., 2017)	82.7
Understanding Conv (Wang et al., 2017)	83.1
RefineNet (Lin et al., 2017a)	82.4
RefineNet-ResNet152 (Lin et al., 2017a)	83.4
PSPNet Zhao et al. (2017)	85.4
SUNet-7-128 (OS = 16)	84.3

Table 6: Performance comparison on PASCAL VOC 2012 test set. For fair comparison, only the methods pre-trained using MS-COCO are displayed. PSPNet uses 82% more parameters.

Methods	mIoU
LRR (VGG16) (Ghiasi & Fowlkes, 2016)	69.7
DeepLabv2+CRF (Chen et al., 2016a)	70.4
Deep Layer Cascade* (Li et al., 2017)	71.1
Piecewise (VGG16) (Lin et al., 2016)	71.6
RefineNet (Lin et al., 2017a)	73.6
Understanding Conv Wang et al. (2017)	77.6
PSPNet Zhao et al. (2017)	78.4
SUNet-7-128 (OS = 16)	75.3
SUNet-7-128 (OS = 8)	76.8

Table 7: Performance comparison on Cityscapes *test* set. All methods were trained only using the “fine” set. All nets utilize ResNet-101 as a base network, except if specified or marked with *.

6.2.2 RESULTS ON TEST SET

PASCAL VOC 2012: Before submitting test set output to an evaluation server, the above model was further fine-tuned on the “trainval” set with batch-norm parameters frozen and at $10\times$ smaller initial learning rate. Table 6 compares the test set results against other state-of-the-art methods. PSPNet performs slightly better than SUNet, but at the cost of 29.1M more parameters while training at an *output_stride* = 8 with $2\times$ runtime. Specifically, the segmentation model of SUNet uses 35.4M parameters while PSPNet requires 64.5M - an increase of 82.2%. The qualitative results can be found in Appendix C.

Cityscapes: A similar training strategy as in PASCAL is adopted except that the multi-scale inference is performed on additional scales $\{1.5, 1.75, 2.0, 2.25, 2.5\}$. Only the *finer* annotation set was used for training. The comparison on the Cityscapes test set results are displayed in Table 7.

7 DISCUSSION AND CONCLUSION

The fundamental structure of conventional bottom-up classification networks limits their efficacy on secondary tasks involving pixel-level localization or classification. To overcome this drawback, a new network architecture, stacked u-nets (SUNets), is discussed in this paper. SUNets leverage the information globalization power of u-nets in a deeper network architecture that is capable of handling the complexity of natural images. SUNets performs exceptionally well on semantic segmentation tasks while achieving fair performance on ImageNet classification. It generalizes well across tasks without any bells and whistles (difference between networks for two task is dilation).

REFERENCES

- Yoshua Bengio, Patrice Simard, and Paolo Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on neural networks*, 5(2):157–166, 1994.
- Achi Brandt. Multi-level adaptive solutions to boundary-value problems. *Mathematics of computation*, 31(138):333–390, 1977.
- William L Briggs, Steve F McCormick, et al. *A multigrid tutorial*, volume 72. Siam, 2000.
- Siddhartha Chandra and Iasonas Kokkinos. Fast, exact and multi-scale inference for semantic image segmentation with deep gaussian crfs. In *European Conference on Computer Vision*, pp. 402–418. Springer, 2016.
- Siddhartha Chandra, Nicolas Usunier, and Iasonas Kokkinos. Dense and low-rank gaussian crfs using deep embeddings. In *ICCV 2017-International Conference on Computer Vision*, 2017.
- 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. *arXiv preprint arXiv:1606.00915*, 2016a.
- Liang-Chieh Chen, Yi Yang, Jiang Wang, Wei Xu, and Alan L Yuille. Attention to scale: Scale-aware semantic image segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3640–3649, 2016b.
- Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. *arXiv preprint arXiv:1802.02611*, 2018.
- Francois Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1251–1258, 2017.
- 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3213–3223, 2016.
- Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 764–773, 2017.
- Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *International journal of computer vision*, 111(1):98–136, 2015.
- Jun Fu, Jing Liu, Yuhang Wang, and Hanqing Lu. Stacked deconvolutional network for semantic segmentation. *arXiv preprint arXiv:1708.04943*, 2017.
- Golnaz Ghiasi and Charless C Fowlkes. Laplacian pyramid reconstruction and refinement for semantic segmentation. In *European Conference on Computer Vision*, pp. 519–534. Springer, 2016.
- Bharath Hariharan, Pablo Arbelaez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *International Conference on Computer Vision (ICCV)*, 2011.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pp. 1026–1034, 2015.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016a.

- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, pp. 630–645. Springer, 2016b.
- Gao Huang, Zhuang Liu, Kilian Q Weinberger, and Laurens van der Maaten. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, volume 1, pp. 3, 2017.
- Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *CVPR*, 2017.
- Simon Jégou, Michal Drozdal, David Vazquez, Adriana Romero, and Yoshua Bengio. The one hundred layers tiramisu: Fully convolutional densenets for semantic segmentation. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*, pp. 1175–1183. IEEE, 2017.
- Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *Advances in neural information processing systems*, pp. 109–117, 2011.
- Xiaoxiao Li, Ziwei Liu, Ping Luo, Chen Change Loy, and Xiaoou Tang. Not all pixels are equal: Difficulty-aware semantic segmentation via deep layer cascade. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3193–3202, 2017.
- Chen Liang-Chieh, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. In *International Conference on Learning Representations*, 2015.
- Guosheng Lin, Chunhua Shen, Anton Van Den Hengel, and Ian Reid. Efficient piecewise training of deep structured models for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3194–3203, 2016.
- Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017a.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pp. 740–755. Springer, 2014.
- Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, volume 1, pp. 4, 2017b.
- Wei Liu, Andrew Rabinovich, and Alexander C Berg. Parsenet: Looking wider to see better. *arXiv preprint arXiv:1506.04579*, 2015.
- Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3431–3440, 2015.
- Ilya Loshchilov and Frank Hutter. Sgdr: stochastic gradient descent with restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *3D Vision (3DV), 2016 Fourth International Conference on*, pp. 565–571. IEEE, 2016.
- Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European Conference on Computer Vision*, pp. 483–499. Springer, 2016.
- Chao Peng, Xiangyu Zhang, Gang Yu, Guiming Luo, and Jian Sun. Large kernel matters—improve semantic segmentation by global convolutional network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4353–4361, 2017.
- Gernot Riegler, Ali Osman Ulusoy, Horst Bischof, and Andreas Geiger. Octnetfusion: Learning depth fusion from data. In *Proceedings of the International Conference on 3D Vision*, 2017.

- 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*, pp. 234–241. Springer, 2015.
- Alexander G Schwing and Raquel Urtasun. Fully connected deep structured networks. *arXiv preprint arXiv:1503.02351*, 2015.
- Abhinav Shrivastava, Rahul Sukthankar, Jitendra Malik, and Abhinav Gupta. Beyond skip connections: Top-down modulation for object detection. *arXiv preprint arXiv:1612.06851*, 2016.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Bharat Singh and Larry S Davis. An analysis of scale invariance in object detection-snip. *arXiv preprint arXiv:1711.08189*, 2017.
- Panqu Wang, Pengfei Chen, Ye Yuan, Ding Liu, Zehua Huang, Xiaodi Hou, and Garrison Cottrell. Understanding convolution for semantic segmentation. *arXiv preprint arXiv:1702.08502*, 2017.
- Zifeng Wu, Chunhua Shen, and Anton van den Hengel. Wider or deeper: Revisiting the resnet model for visual recognition. *arXiv preprint arXiv:1611.10080*, 2016.
- Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015.
- Fisher Yu, Vladlen Koltun, and Thomas Funkhouser. Dilated residual networks. In *Computer Vision and Pattern Recognition*, volume 1, 2017.
- Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp. 2881–2890, 2017.
- Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip HS Torr. Conditional random fields as recurrent neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1529–1537, 2015.

APPENDIX

A IMAGENET: IMPLEMENTATION DETAILS

All the models were implemented using the PyTorch deep learning framework and trained using four P6000 GPUs on a single node. We use SGD with a batch size of 256. For our largest model, 7-128, we were limited to a batch size of 212, due to the GPUs memory constraints. The initial learning rate was set to 0.01 and decreased by a factor of 10 every 30 epochs. We use a weight decay of $5e^{-4}$ and Nesterov momentum of 0.9 without dampening. The weights were initialized as in He et al. (2015) and all the models were trained from scratch for a total of 100 epochs.

B SEMANTIC SEGMENTATION

A DATASETS

PASCAL VOC 2012: This dataset contains 1,464 train, 1,449 validation and 1,456 test images. The pixel-level annotation for 20 objects and one background class is made available for the train and validation set. Following common practice, the train set is augmented with additional annotated data from Hariharan et al. (2011) which finally provides a total of 10,582 (trainaug) training images.

Cityscape: This dataset consists of finely annotated images of urban scenes covering multiple instances of cars, roads, pedestrians, buildings, etc. In total, it contains 19 classes on 2,975 finely annotated training and 500 validation images.

B IMPLEMENTATION DETAILS

We use the SGD optimizer with a momentum of 0.95 and weight decay of 10^{-4} . Each model is fine-tuned starting with an initial learning rate of 0.0002 which is decreased every iteration by a factor of $0.5 \times \left(1 + \cos\left(\pi \frac{\text{iter}}{\text{max iters}}\right)\right)$ Loshchilov & Hutter (2016). The batch-norm parameters are learned with a decay rate of 0.99 and the input crop size for each training image is set to 512×512 . We train each model using two P6000 GPUs and the batch size 22. On PASCAL VOC, each model is trained for 45K iterations while for Cityscapes we use 90K iterations.

Data Augmentation: To prevent overfitting during the training process, each training image is resized with a random scale from 0.5 to 2 following which the input image is randomly cropped. Additionally, the input is randomly flipped horizontally and also randomly rotated between -10° to 10° .

C QUALITATIVE RESULTS

The segmentation results for PASCAL VOC 2012 val and test images are displayed in figure 3 and 4 respectively.

D ACTIVATION MAPS

Figure 5 shows the activation map recorded at the end of each level (as indicated in figure 2) for an example input image of an ‘‘Aeroplane.’’ As noted earlier, the inclusion of strided convolutions instead of multigrid dilations leads to noisy feature maps (see col 3; rows 4-6). The addition of de-gridding layers serves to produce a coherent prediction map at the output (see col 2; row 6).




Figure 5: Activation map recorded at the end of each level of the dilated SUNet for an example input image of an ‘Aeroplane’. The activation map with total highest magnitude were selected from among all feature map outputs at the corresponding layer. **Top to Bottom:** output at end of level 1 – 6 followed by classification output. The level 6 output is simply a prediction map before bilinear interpolation.