

Design of Real-time Semantic Segmentation Decoder for Automated Driving

Arindam Das¹, Saranya Kandan¹, Senthil Yogamani² and Pavel Křížek³

¹*Detection Vision Systems, Valeo India, India*

²*Valeo Vision Systems, Valeo Ireland, Ireland*

³*Valeo R&D DVS, Prague, Czech Republic*

Keywords: Semantic Segmentation, Visual Perception, Efficient Networks, Automated Driving.

Abstract: Semantic segmentation remains a computationally intensive algorithm for embedded deployment even with the rapid growth of computation power. Thus efficient network design is a critical aspect especially for applications like automated driving which requires real-time performance. Recently, there has been a lot of research on designing efficient encoders that are mostly task agnostic. Unlike image classification and bounding box object detection tasks, decoders are computationally expensive as well for semantic segmentation task. In this work, we focus on efficient design of the segmentation decoder and assume that an efficient encoder is already designed to provide shared features for a multi-task learning system. We design a novel efficient non-bottleneck layer and a family of decoders which fit into a small run-time budget using VGG10 as efficient encoder. We demonstrate in our dataset that experimentation with various design choices led to an improvement of 10% from a baseline performance.

1 INTRODUCTION

Semantic segmentation provides complete semantic scene understanding wherein each pixel in an image is assigned a class label. It has applications in various fields including automated driving, augmented reality and medical image processing. The complexity of Convolution Neural Networks (CNN) architectures have been growing consistently. However for industrial applications, there is a computational bound because of limited resources on embedded platforms. It is essential to design efficient models which fit the run-time budget of the system. There are many papers which demonstrate large run-time improvements with minimal loss of accuracy by using various techniques. An overview of efficient CNN for semantic segmentation is provided in (Briot et al., 2018).

Semantic segmentation architectures typically have an encoder and a decoder as shown in Figure 1. The encoder extracts features from the image which is then decoded to produce semantic segmentation output. ImageNet pre-trained networks are typically used as encoder. In early architectures (Badrinarayanan et al., 2015) (Ronneberger et al., 2015), decoder was a mirror image of encoder and had the same complexity. Newer architectures use a relatively smaller de-

coder. There can also be additional connections from encoder to decoder. For example, SegNet (Badrinarayanan et al., 2015) passes max-pooling indices and U-net (Ronneberger et al., 2015) passes intermediate feature maps to decoder.

In this paper, we propose the design of a novel non-bottleneck layer particularly to perform semantic segmentation task where the encoder is task independent unlike existing methods. Our non-bottleneck layer based on residual learning, has been designed to perform well for some classes that are not well represented in the dataset. Having cascaded skip connections make our non-bottleneck layer capable to handle high gradient flow and suitable for an embedded platform to run on real time with lightweight encoder model. Table 1 presents the efficacy of the proposed decoder architecture over many variants of the decoder model while maintaining the encoder to be task independent and constant.

The rest of the paper is structured as follows. Section 2 discusses the related work on semantic segmentation. 3 explains the details of the proposed network. Section 4 details the experimental setup and results. Finally, section 5 summarizes the paper and provides potential future directions.

2 BACKGROUND

The main purpose of convolutional neural network was to perform image classification so that an object of a particular category can be predicted irrespective of its rotation, translation and scale. But Long *et al.* (Long *et al.*, 2015) considered the same concept as encoder and added fully convolutional upsampling layer using the concept of unpooling to construct the output. Later deconvolution technique was introduced instead of unpooling in (Badrinarayanan *et al.*, 2015) along with considering skip connections to pull encoded feature maps. As part of the recent advancement in semantic segmentation tasks, high usage of dilated convolution is observed in many works (Chen *et al.*, 2018a) (Chen *et al.*, 2018b) (Hamaguchi *et al.*, 2017) (Romera *et al.*, 2017) (Wang *et al.*, 2018) (Yu and Koltun, 2015) (Zhao *et al.*, 2017) (Yang *et al.*, 2018), this is because it helps to grow the receptive field exponentially without loss of resolution. As notable other works, DeepLab2 (Chen *et al.*, 2018a) uses spatial pyramid pooling in ResNet-101 and it includes Conditional Random Fields (CRF) as well. RefineNet (Lin *et al.*, 2017) proposed a multi-path refinement network that exploits information to perform better segmentation in higher resolution. It basically refines features from low resolution in recursive manner to obtain better prediction in high resolution. Though none of the methods use task independent lightweight encoder model as used in this work. Pohlen *et al.* (Pohlen *et al.*, 2017) proposed a ResNet like architecture that considers multi-scale input in two stream process. One stream extracts features from full resolution to get knowledge about semantic boundaries and the other one works on different dimension as received from sequence of pooling layers to extract robust local features. ENet (Paszke *et al.*, 2016) presented a network aimed for mobile devices where ResNet like architecture is considered as encoder and it contains dilated convolution to extract semantic features. Another realtime segmentation work, ICNet (Zhao *et al.*, 2017) reported a multi-branch architecture where inputs in different resolution are considered and it developed cascade feature fusion unit across multiple branches. In one more recent paper, ContextNet (Poudel *et al.*, 2018) proposed multi-column encoder architecture and later fusion at decoder. A deep network as encoder is used to extract contextual features from relatively smaller resolution and another encoder that is shallow in nature is added to extract global features.

The methods mentioned above do not use task independent lightweight encoder model as used in this work and do not support multi task learning practice.

3 PROPOSED METHOD

3.1 Overview

The reported work is conceptualized for the embedded platform where the entire system can run on real time, thus priority is execution speed over accuracy. Similar to existing neural network architectures such as FCN (Long *et al.*, 2015), SegNet (Badrinarayanan *et al.*, 2015), UNet (Ronneberger *et al.*, 2015) for segmentation, the proposed network has also two parts which are encoder and decoder, which is furnished in Figure 1. An encoder is basically a Convolutional Neural Network (CNN) that helps to extract features from various feature dimensions and thus reduces the problem space as it becomes deeper. A decoder does the opposite of encoder, it consumes the feature maps from last layer as well as intermediate layers of encoder and reconstructs the original input space. For reconstruction we use deconvolution layer instead max-unpooling as used in SegNet (Badrinarayanan *et al.*, 2015) and ENet (Paszke *et al.*, 2016). Once the reconstruction is done then the decoder generates class prediction for each pixel in original input space.

3.2 Encoder

It is mentioned earlier that the intention of this study is to develop a lightweight decoder for segmentation task where the encoder is task independent. It means that the features extracted by the single encoder will be shared across all decoders to accomplish separate tasks (segmentation, detection, classification etc.). In the context of our work, it is observed that the encoder design is mostly specific to the semantic segmentation task such as (Badrinarayanan *et al.*, 2015), (Romera *et al.*, 2017), (Treml *et al.*, 2016), (Yu and Koltun, 2015), thus these encoders had all necessary components to extract semantic features and the decoder performs well with the readily available information. However, our study is intended towards use of a more generic encoder. So we designed a task independent, designed a VGG (Simonyan and Zisserman, 2014) style classifier of 10 layers as encoder and no special attention given for semantic segmentation such as heavy usage of dilated convolution as used in (Chen *et al.*, 2018a) (Chen *et al.*, 2018b) (Hamaguchi *et al.*, 2017) (Romera *et al.*, 2017) (Wang *et al.*, 2018) (Yu and Koltun, 2015) (Zhao *et al.*, 2017) (Yang *et al.*, 2018) or information fusion at different resolution (Lin *et al.*, 2017) (Poudel *et al.*, 2018) (Zhao *et al.*, 2017) (Zhang *et al.*, 2018).

As per Figure 3, convolution with stride 2 followed by max-pooling is used to reduce the problem

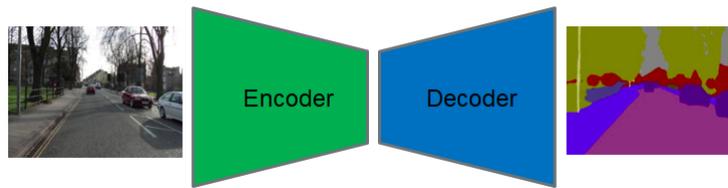


Figure 1: Typical encoder-decoder architecture of CNN based semantic segmentation network.

space, thus reducing the number of hyperparameters and also the run-time. Obviously, this is a trade-off for segmentation accuracy, but that is not the case for other tasks like detection, classification etc. Considering this encoder to be function independent, this gap of spatial information exploration needs to be overcome in the decoder side by learning semantic features extensively. Convolution layers are added sequentially along with increasing width and decreasing feature space in regular interval. All the convolution layers use kernel of size 5×5 followed by Batch-Normalization (Ioffe and Szegedy, 2015) to speed up the convergence and ReLU (Nair and Hinton, 2010) to add non-linearity. The last layer of encoder produces total 256 feature maps.

3.3 Decoder

The proposed architecture of decoder has been obtained after redesigning the existing features of the convolutional nets, such as residual learning (He et al., 2016) and non-bottleneck layers (Romera et al., 2017). In the recent past, learning through residual blocks has shown breakthrough performance in many vision related tasks and made it possible to make the network grow more deeper very easily. This advancement helped to outperform significantly in many object recognition tasks. Further the same learning strategy has been also used for semantic segmentation as in (Romera et al., 2017). However, in (Das and Yogamani, 2018), it has been demonstrated that residual learning for a network with lesser depth is not efficient. This is because the network can not handle high gradient flow during back-propagation. To circumvent this issue, in this study, the original residual learning (He et al., 2016) strategy has been modified with an adequate arrangement to distribute the high gradients through multiple skip connections.

As per our design of the encoder, it can be realized that there has been no suitable mechanism employed to extract semantic features. Now, following with the recent trend, if the decoder is going to have only few set of deconvolution layers to reconstruct the output then the segmentation result will be definitely poor. The reason is the performance of the decoder is limited to the knowledge shared by the encoder and as

per our design, the encoder is expected to share knowledge that is non-semantic in nature. To address this affair, there is a requisite to learn semantic information from the encoded feature space in the decoder. Hence, we need non-bottleneck blocks between two deconvolution layers. In (Romera et al., 2017), the authors have designed non-bottleneck layer which is 1D in nature and claimed to be a better regularizer and also faster. The same non-bottleneck layer is used in the encoder as well and most of the convolutions are dilated.

The design of our non-bottleneck layer is shown in Figure 2. It contains both 1D and 3D kernels. 1D kernel is used to extract information mainly in one direction at a time and 3D kernel is for gathering features from bigger receptive area. Later we try to look for dense information through multiple kernels with different sizes for example 3×3 , 5×5 and 1×1 respectively. Following this, the features extracted using different kernels are fused. This method helps to summarize the semantic features that are collected from different receptive area. The resultant features are again fused with the input features to the same non-bottleneck layer. The multiple skip connections to the feature fusion blocks in the proposed non-bottleneck layer help to handle high gradient flow because during back-propagation the incoming gradient gets distributed among all paths. As per Figure 2, our non-bottleneck layer has two variants that are type-1 (left) and type-2 (right). The only difference between two variants is the block at the right uses dilated convolution where the kernel size is 3×3 and 5×5 . Dilated convolution helps to extract spatial information by expanding the field-of-view as per the dilation factor while maintaining the same resolution. With increasing dilation rate, receptive field is also expanded however, for the present work the dilation rate has been kept constant and it is 1. It is possible to receive better accuracy with incremental dilation rate but dilated convolution is computationally expensive and it becomes more costlier as the dilation rate increases. Considering the present study to be aimed for embedded platform, usage of similar features need to be employed carefully to meet the run-time. After each convolution operation, ReLU (Nair and Hinton, 2010) is used as activation unit for better convergence

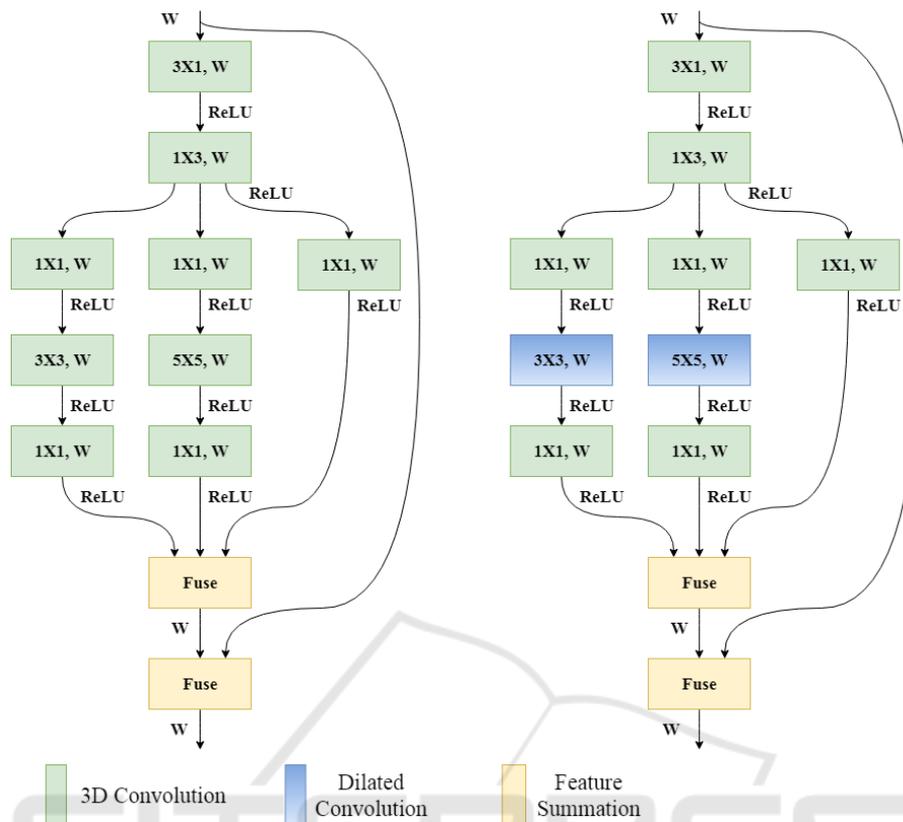


Figure 2: Illustration of non-bottleneck layers.

(Krizhevsky et al., 2012).

3.4 Encoder-Decoder Architecture

Figure 3 shows the encoder-decoder pair that is used in this work. The encoder is a very generic one and the decoder is our main proposal. The feature maps passed from the encoder are downsampled from 256 to 4 as the present study concentrates on 4 classes. One can argue that drastic change in the number of feature maps such as this can impact on accuracy but if we make the decoder wider that will shoot up the runtime significantly. Thus decreasing the number of feature maps in regular interval is not affordable and out of our budget. After first deconvolution layer, non-bottleneck layer of type-1 is used two times. Following second and third deconvolution layers, non-bottleneck layers of type-2 are used twice and once respectively. There is no non-bottleneck layer used after fourth deconvolution. The non-bottleneck layers are designed in such a way that first repetitive usage of type-1 as intermediate block between two deconvolution layers and then having multiple type-2 especially between the later deconvolution layers will help to remove the gap of learning semantic information that

we have seen in the encoder side. Hence, our proposed non-bottleneck layer is generic and can be used for any segmentation task. Also we pulled intermediate feature maps from the second and third last convolution layers of the encoder. This has to be highlighted that our encoder consumed most of the run time because it is much wider than the decoder. To circumvent this issue, we could use the concept of group convolution as proposed by Xie et al. (Xie et al., 2017) but in (Das et al., 2018), it has been experimentally demonstrated that aggregated residual transformation shows adverse effect for lightweight networks such as our encoder.

Details on how to train our proposed network end-to-end, other training strategy and hyperparameter details, are discussed in the next section.

4 EXPERIMENTS

We have performed thorough investigation to show the efficacy and robustness of our proposed network.



Figure 3: CNN based encoder-decoder architecture used for semantic segmentation task.

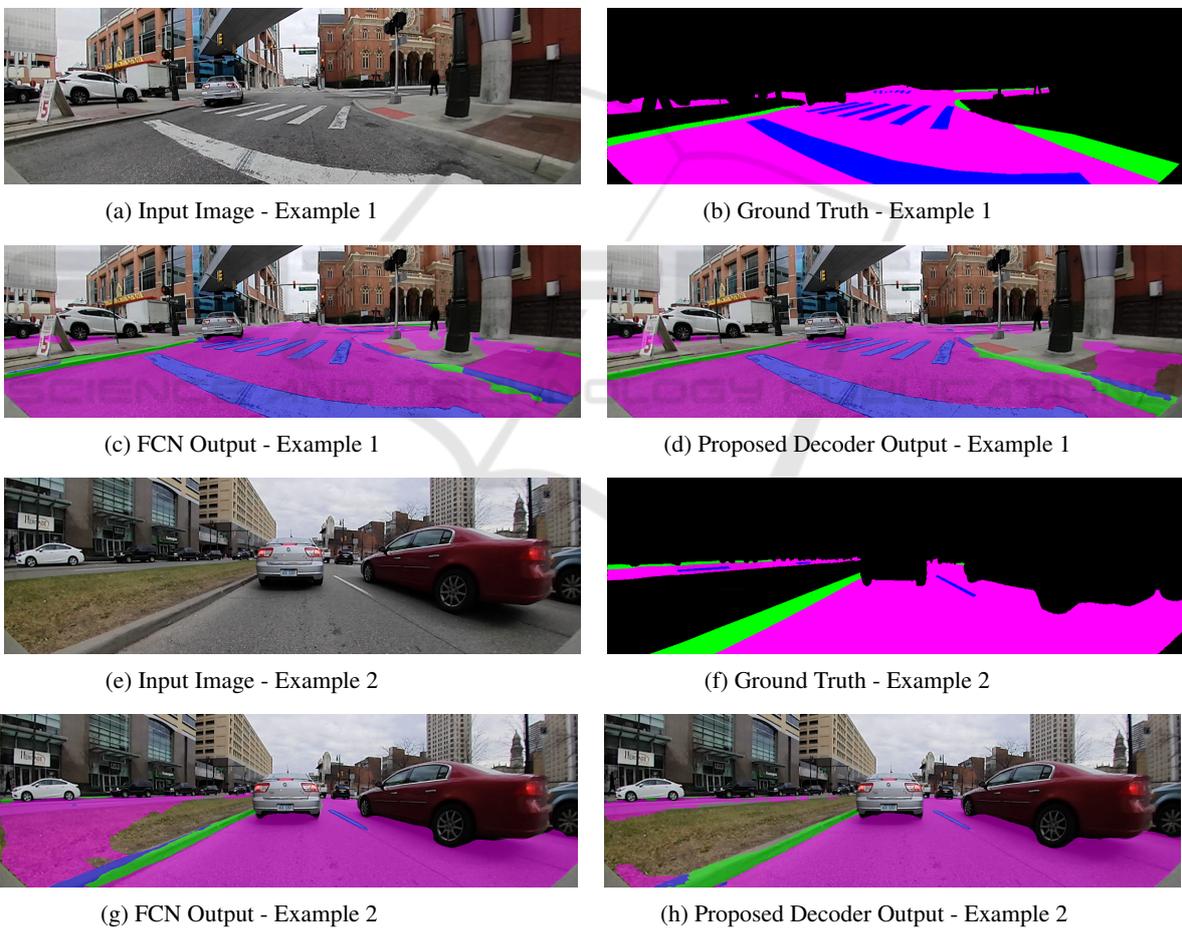


Figure 4: Comparison of proposed optimal decoder against a standard FCN decoder.

Dataset. We have used dataset that is owned by our organization. However, some samples of the dataset are shown in Figure 4. In this dataset, dimension of each image is 1280X384. Samples are highly va-

ried in nature and mostly urban scenes. Diverse lighting conditions and presence of dense shadow make this dataset challenging. In the present study, to perform semantic segmentation, we concentrated only

on 4 critical classes that are lanes, curb, road and void (everything else). The entire dataset is divided into Training, Validation and Test set each containing 3016, 981 and 1002 images respectively. Corresponding each sample image, single channel annotation was developed where each pixel has only one class label.

Training. First we train our encoder from scratch on ImageNet and then transfer the weights for the present task. This transfer learning will help to get better initialization of the weights at the beginning of the training for semantic segmentation. The pre-training on a much larger dataset was required because our model is quite shallow and the dataset used in this work is very small in size. Consequently there is a high chance that the network will suffer from overfitting due to lack of training samples. We could follow the concept of layer-wise training in a supervised fashion to train the encoder to extract more robust features as reported in (Roy et al., 2016) as the encoder used in this work is quite lightweight. Implementation of the proposed network and all experiments are executed using Keras (Chollet et al., 2015) framework. We considered very popular Adam (Kingma and Ba, 2014) as optimizer. Regarding the other network configuration, weight decay and batch size were set to 0.9 and 4 respectively. Training was started with 0.0005 as initial learning rate including standard polynomial decay strategy to decrease this value over 350K iterations. Dropout (Srivastava et al., 2014) is not used in our model. For all experiments, we used NVIDIA Titan X 12G GPU with 24 GB RAM. No data augmentation technique has been performed during training.

Experimental Results and Comparison Study.

The hardware that we use is designed with automotive power constraints in mind, thus having restricted number of features for design of a convolutional neural network. Also we intend to utilize a generic encoder network that stays well within the budget. Thus our main objective is to design an efficient decoder that satisfies both these constraints. In the course of design of an efficient decoder, we have experimented multiple versions of decoder all of which are explained later in this section. With all these decoders VGG10 pre-trained with Imagenet is used as the encoder. This will help to have fair comparison of the different variants of the decoder. The entire network containing the encoder along with different decoders is trained end to end with the available pixel-wise ground truth label. All these variants of decoder reported in Table 1, fits all our constraints. It is quite general while capturing an urban scene, there will be

very limited region occupied by lanes, curbs but it is exactly opposite for roads and void classes. So we clearly see that for effective learning there is a huge gap in problem space in these two classes while comparing with other. Without even attempting any data augmentation and class weighing technique, our non-bottleneck layers worked better for curb though there is a slight deterioration for lanes. To evaluate the segmentation performance on all the designed decoders, widely used Intersection over Union (IoU) metric is considered and details are furnished in Table 1.

As put forth earlier, we did experiments with several combinations of Non-Bottleneck layers in the network whose results are updated in Table 1. Decoder D1 uses our proposed non-bottleneck layer without 1X1 convolution after 3X3 and 5X5 convolution. Decoder D2 is same as D1 but it does not use second skip connection from encoder. Decoder D3 shares the same configuration as D2 but the batch size during training was 8 whereas it was 4 for D2. Decoder D4 is same as D3 but it does not use 1X1 convolution even before 3X3 and 5X5 convolution. Decoder D5 is a bit different. After first deconvolution layer two sets of 3X1, 1X3 convolutions followed by ReLU is used. Also it uses skip connection to fuse the resultant features with the input feature maps of first 3X1 convolution. After second deconvolution layer, one 3X3 dilated convolution with dilation rate 1 is used and then the same non-bottleneck as used after first deconvolution layer. Only 3X3 dilated convolution with dilation rate 1 is used after third and fourth deconvolution layer. Decoder D6 is same as D5 but it uses batch size as 4 where 8 was used in D5. Decoder D7 is different in terms of kernel size in deconvolution layers. It uses kernel of size 2X2 in first and second, 3X3 in third and fourth, 5X5 in fifth upsampling layer. Decoder D8 uses same non-bottleneck as D7 without 3X3 and 5X5 convolution. In the pattern mNp in Table 1, m stands for number of non-Bottleneck (N) layers, p stands for the type of non-bottleneck layer. Representation within braces (and) stands for set of non-bottleneck layers after a deconvolution layer starting from the first one. Further to explore more, we have modified the design of our non-bottleneck layer from various aspects to check what variant of change seems to work better. Please note that the modified non-bottleneck layers are non-repetitive between two deconvolution layers. Of the different decoder variant, the best version is the one put forth in Figure 3, which is obtained after several optimization efforts and this network uses the Non-bottleneck layers detailed in Figure 2. This network also takes care of the class imbalance for Lanes and Curb and improves its class-wise IoU. We also compared the runtime

Table 1: Results of different variants of our decoder with non-bottleneck layer and VGG10 as encoder for semantic segmentation on our dataset.

Decoder configuration	Avg. class accuracy			Avg. class IoU score			Mean
	Lanes	Curb	Road	Lanes	Curb	Road	
D1	0.5678	0.6228	0.9688	0.4943	0.4296	0.9265	0.7255
D2	0.4363	0.6082	0.9694	0.4021	0.4445	0.9198	0.6989
D3	0.5096	0.5682	0.9689	0.4333	0.4264	0.9254	0.7038
D4	0.4769	0.5932	0.9684	0.4152	0.4534	0.9248	0.7013
D5	0.4561	0.6059	0.9668	0.4013	0.403	0.9138	0.6866
D6	0.5775	0.6539	0.9566	0.4567	0.4198	0.9154	0.7106
D7	0.5263	0.4284	0.9464	0.4192	0.2911	0.8689	0.6458
D8	0.5628	0.6755	0.9686	0.4752	0.4717	0.9296	0.727
(2N1)(2N1)(2N2)(2N2)	0.6054	0.6534	0.9685	0.5294	0.4755	0.9299	0.742
(2N2)(2N2)(2N2)(2N2)	0.5951	0.647	0.9655	0.5061	0.5061	0.9267	0.7323
(2N2)(2N2)(2N2)	0.6	0.639	0.9661	0.5063	0.4391	0.9279	0.7315
(2N1)(2N1)(2N1)(2N1)	0.5743	0.6343	0.9691	0.4961	0.4486	0.9275	0.7284
(1N1-1N2)(1N1-1N2)(1N1-1N2)(1N1-1N2)	0.5938	0.6338	0.9681	0.5051	0.4438	0.9264	0.731
Optimal	0.6118	0.6588	0.9689	0.5304	0.4696	0.9314	0.7441

measurement of this network with the popular FCN network on our custom hardware. We find that this network has the advantage of improving the IoUs especially of the key classes with less than 2ms increase in runtime. The sample segmentation outputs of our proposed optimal decoder and a standard FCN decoder are shown in Figure 4.

5 CONCLUSION

Design of efficient encoders is a growing area of research. In this work, we focused on design of efficient decoders. Firstly, we designed a novel efficient non-bottleneck layer and a family of decoders based on this layer. We experimentally show that different choice of decoder design had a large impact and the optimal configuration had 10% improvement of accuracy in terms of mean IoU over a baseline configuration. In particular, for more difficult segmentation classes like lanes and curb, higher improvements of 12% and 18% were obtained. Thus we demonstrate that the design of an efficient decoder can play a critical role for segmentation tasks as it covers a significant portion of the overall computation of the network. We hope that our preliminary study demonstrates the need for further research on efficient decoders. In future work, we build a systematic family of meta-architectures with a fixed run-time budget and learn the optimal configuration using meta-learning techniques.

REFERENCES

- Badrinarayanan, V., Kendall, A., and Cipolla, R. (2015). Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *arXiv preprint arXiv:1511.00561*.
- Briot, A., AI, G., Creteil, V., Viswanath, P., and Yogamani, S. (2018). Analysis of efficient cnn design techniques for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 663–672.
- Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., and Yuille, A. L. (2018a). Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848.
- Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., and Adam, H. (2018b). Encoder-decoder with atrous separable convolution for semantic image segmentation. *arXiv preprint arXiv:1802.02611*.
- Chollet, F. et al. (2015). Keras. <https://keras.io>.
- Das, A., Boulay, T., Yogamani, S., and Ou, S. (2018). Evaluation of group convolution in lightweight deep networks for object classification. In *Proceedings of the 24th International Conference on Pattern Recognition Workshop*.
- Das, A. and Yogamani, S. (2018). Evaluation of residual learning in lightweight deep networks for object classification. In *Proceedings of the 20th Irish Machine Vision and Image Processing Conference*, pages 205–208.
- Hamaguchi, R., Fujita, A., Nemoto, K., Imaizumi, T., and Hikosaka, S. (2017). Effective use of dilated convolutions for segmenting small object instances in remote sensing imagery. *CoRR*, abs/1709.00179.

- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Ioffe, S. and Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105.
- Lin, G., Milan, A., Shen, C., and Reid, I. D. (2017). Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Cvpr*, volume 1, page 5.
- Long, J., Shelhamer, E., and Darrell, T. (2015). Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440.
- Nair, V. and Hinton, G. E. (2010). Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 807–814.
- Paszke, A., Chaurasia, A., Kim, S., and Culurciello, E. (2016). Enet: A deep neural network architecture for real-time semantic segmentation. *arXiv preprint arXiv:1606.02147*.
- Pohlen, T., Hermans, A., Mathias, M., and Leibe, B. (2017). Fullresolution residual networks for semantic segmentation in street scenes. *arXiv preprint*.
- Poudel, R. P., Bonde, U., Liwicki, S., and Zach, C. (2018). Contextnet: Exploring context and detail for semantic segmentation in real-time. *arXiv preprint arXiv:1805.04554*.
- Romera, E., Alvarez, J. M., Bergasa, L. M., and Arroyo, R. (2017). Efficient convnet for real-time semantic segmentation. In *IEEE Intelligent Vehicles Symp.(IV)*, pages 1789–1794.
- Ronneberger, O., Fischer, P., and Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer.
- Roy, S., Das, A., and Bhattacharya, U. (2016). Generalized stacking of layerwise-trained deep convolutional neural networks for document image classification. In *Pattern Recognition (ICPR), 2016 23rd International Conference on*, pages 1273–1278. IEEE.
- Simonyan, K. and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958.
- Treml, M., Arjona-Medina, J., Unterthiner, T., Durgesh, R., Friedmann, F., Schuberth, P., Mayr, A., Heusel, M., Hofmarcher, M., Widrich, M., et al. (2016). Speeding up semantic segmentation for autonomous driving. In *MLITS, NIPS Workshop*.
- Wang, P., Chen, P., Yuan, Y., Liu, D., Huang, Z., Hou, X., and Cottrell, G. (2018). Understanding convolution for semantic segmentation. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1451–1460. IEEE.
- Xie, S., Girshick, R., Dollár, P., Tu, Z., and He, K. (2017). Aggregated residual transformations for deep neural networks. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 5987–5995. IEEE.
- Yang, M., Yu, K., Zhang, C., Li, Z., and Yang, K. (2018). Denseaspp for semantic segmentation in street scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3684–3692.
- Yu, F. and Koltun, V. (2015). Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*.
- Zhang, Z., Zhang, X., Peng, C., Cheng, D., and Sun, J. (2018). Exfuse: Enhancing feature fusion for semantic segmentation. *arXiv preprint arXiv:1804.03821*.
- Zhao, H., Qi, X., Shen, X., Shi, J., and Jia, J. (2017). Icnnet for real-time semantic segmentation on high-resolution images. *arXiv preprint arXiv:1704.08545*.