

Nr. 59  
18. January 2019

Preprint-Series: Department of Mathematics - Applied Mathematics

Projection-Based 2.5D U-net Architecture for  
Fast Volumetric Segmentation

C. Angermann, M. Haltmeier, R. Steiger, S. Pereverzyev and  
E. Gizewski



---

Technikerstraße 13 - 6020 Innsbruck - Austria  
Tel.: +43 512 507 53803 Fax: +43 512 507 53898  
<https://applied-math.uibk.ac.at>

# Projection-Based 2.5D U-net Architecture for Fast Volumetric Segmentation

Christoph Angermann and Markus Haltmeier

Department of Mathematics  
University of Innsbruck  
Technikerstrasse 13, 6020 Innsbruck, Austria  
{christoph.angermann,markus.haltmeier}@uibk.ac.at

Ruth Steiger, Sergiy Pereverzyev Jr., and Elke Gizewski

Universitätsklinik für Neuroradiologie  
Medizinische Universität Innsbruck  
Anichststraße 35, 6020 Innsbruck  
neuroradiologie@i-med.ac.at

February 1, 2019

## Abstract

Convolutional neural networks are state-of-the-art for various segmentation tasks. While for 2D images these networks are also computationally efficient, 3D convolutions have huge storage requirements and require long training time. To overcome this issue, we introduce a network structure for volumetric data without 3D convolution layers. The main idea is to integrate projection layers to transform the volumetric data to a sequence of images, where each image contains information of the full data. We then apply 2D convolutions to the projection images followed by lifting to a volumetric data. The proposed network structure can be trained in much less time than any 3D-network and still shows accurate performance for a sparse binary segmentation task.

## 1 Introduction

Deep convolutional neural networks have become a powerful method for image recognition ([7, 4]). In the last few years they also exceeded the state-of-the-art in providing segmentation masks for images. In [5], the idea of using fully convolutional networks for segmentations came up. Based on this work, the U-net introduced in [6] provides a powerful 2D segmentation tool for biomedical applications. It has been demonstrated to learn highly accurate ground truth masks from only very few training samples.

Among others, the fully automated generation of volumetric segmentation masks becomes increasingly important for biomedical applications. This task still is challenging. One idea is to extend the U-net structure to volumetric data by using 3D convolutions, as has been proposed in [9, 2]. Essential drawbacks are the huge memory requirements and long training time. Deep learning segmentation methods therefore are often applied to 2D slice images. However, these slice images do not contain information of the full 3D data which makes the segmentation much more complex.

To address the drawbacks of existing approaches, we introduce a network structure which is able to generate accurate volumetric segmentation masks of arbitrary large 3D volumes. The main idea is to integrate projection layers from different directions which transform the data to 2D images containing information of the full image. As an example we test the network for segmenting blood vessels in magnetic resonance angiography (MRA) scans. The proposed network demonstrates to be nearly as fast as 2D-networks without using sliding-window techniques, requires order of magnitude less memory than networks with 3D convolutions and still produces accurate results.

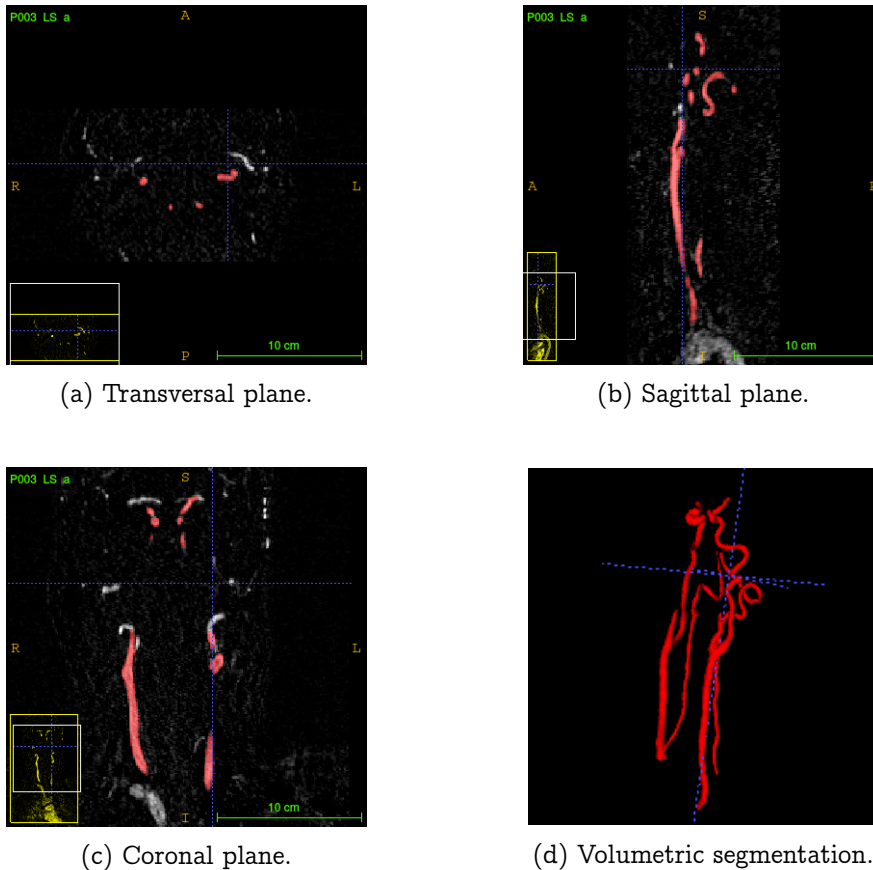


Figure 1.1: In every plane the blood vessels of interest are marked in red. In the fourth picture, we see the resulting segmentation mask. The segmentation was conducted with the freeware ITK-SNAP [8].

## 2 Background

### 2.1 Volumetric segmentation of blood vessels

As our targeted application, we aim at generating volumetric binary segmentation masks. In particular, we aim at segmenting blood vessels (arteries and veins) which assists the doctor to detect abnormalities like stenosis or aneurysms. Furthermore, the medical sector is looking for a fully automated method to evaluate large cohorts in the future. The Department of Neuroradiology Innsbruck has provided volumetric MRA-scans of 119 different patients. The images face the arteries and veins between the brain and the chest. Fortunately, also the volumetric segmentation masks (ground truths) of these 119 patients have been provided. These segmentation masks have been generated by hand which is long hard work. The found segmentation mask for one of these scans is shown in Figure 1.1.

Our goal is the fully automated generation of the volumetric segmentation masks of the blood vessels. For that purpose we use deep learning and neural networks. At the first glance, this problem may seem to be quite easy because we only have two labels (0: background, 1: blood vessel). But we have to notice that we do not want to segment all vessels but only the vessels of interest. So there are also arteries and veins which have label 0 which might confuse the network. Other challenges are caused by the big size of the volumes ( $96 \times 288 \times 224$  voxels) and by the very unbalanced distribution of the two labels (in average 99.76 % of all voxels indicate background in the ground truth).

### 2.2 Segmentation of MIP images

We first solve a 2D version of our problem. This can be done by applying maximum intensity projection (MIP) to the 3D data and the corresponding 3D segmentation masks. Using a rotation angle of  $\alpha = 36^\circ$  around the vertical axis we obtain MIPs out of each patient, which results in a data set to 1190 pairs of 2D images and corresponding 2D segmentation masks. Data corresponding for one of the patients are shown in Figure 2.1.

The U-net used for binary segmentation is a mapping  $\mathcal{U}: \mathbb{R}^{a \times b} \rightarrow [0, 1]^{a \times b}$  which takes an image as input and outputs for each pixel the probability of being a foreground pixel. It is formed by the following ingredients [6]:

- The *contracting part*: It includes stacking over convolution and max-pooling layers considering following properties: (1) We only use  $3 \times 3$  filters to hold down complexity and use zero-padding to guarantee that all layer outputs have even spatial dimension. (2) Each max-pooling layer has stride (2, 2) to half the spatial dimensions. We must be aware that the spatial dimensions of the input image can get divided by 2 often enough without producing any rest. This can be done by slight cropping. (3) After each max-pooling layer we augment number of filters by the factor 2.
- The *upsampling part*: To obtain similarity to the contracting part, we make use of transposed convolutions to double spatial dimension and to halve the




Figure 2.1: MIP images of a 3D MRA scan with  $\alpha = 36^\circ$ . In the first row, we see the projections of the original scan, in the second row we see the corresponding projections of the ground truths.

filter size. They are followed by convolutional blocks consisting of two convolution layers with kernel size  $3 \times 3$  after each upsampling layer.

Every convolution layer in this structure gets followed by a ReLu-activation-function. To link the contracting and the upsampling part, concatenation layers are used, where two images with same spatial dimension get concatenated over their channel dimension (see Figure 2.2). This ensures a combination of each pixel’s information with its localization. At the end the sigmoid-activation-function gets applied to obtain for each pixel the probability for being a foreground pixel. To get the final segmentation mask, a threshold (usually 0.5) is applied point-wise to the output of the U-net.

Our implemented U-net has filter size 32 at the beginning and filter size 512 at the end of the contracting part. It is trained with the Dice-loss function [2]

$$\ell(y, \hat{y}) = 1 - \frac{2 \sum_k (y \odot \hat{y})_k}{\sum_k \hat{y}_k + \sum_k y_k}, \quad (2.1)$$

where  $\odot$  denotes pixelwise multiplication, the sums are taken of pixel locations,  $\hat{y} = \hat{f}(x)$  are the probabilities predicted by the U-net and  $y$  is the corresponding ground truth. The Dice-loss function measures similarity by comparing all correctly predicted vessels pixels with the total number of vessels pixels in the prediction.

For  $i, j \in \{0, 1\}$  let us denote by  $p_{ij}$  the set of all pixels of class  $i$  predicted to class  $j$ , and by  $t_i$  the number of all pixels belonging to class  $i$ . With this notation, we evaluate the following evaluation metrics during training:

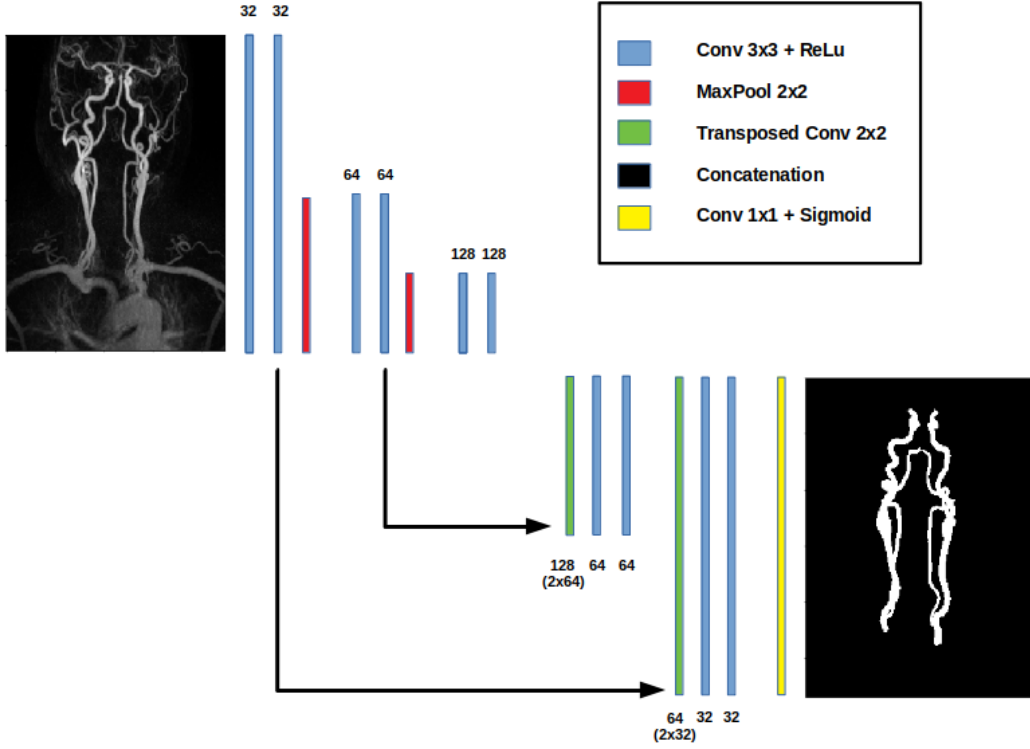


Figure 2.2: Visualization of the ground structure of an U-net for semantic segmentation.

- *Mean Accuracy:*

$$MA \triangleq \frac{1}{2} \left( \frac{p_{00}}{t_0} + \frac{p_{11}}{t_1} \right).$$

- *Mean Region Intersection over Union (IU):*

$$IU \triangleq \frac{1}{2} \left( \frac{p_{00}}{t_0 + p_{10}} + \frac{p_{11}}{t_1 + p_{01}} \right).$$

- *Dice-coefficient:*

$$DC \triangleq \frac{2p_{11}}{2p_{11} + p_{01} + p_{10}}.$$

We also make use of batch normalization layers to speed up convergence and of dropout layers to handle overfitting [3]. For training the Adam-optimizer [1] is used with learning rate 0.001 in combination with learning-rate-scheduling. If the network shows no improvement for 4 epochs the training process gets stopped (early stopping) and the weights of the best epoch in terms of validation loss get restored. Here we use (70, 15, 15) split in training, validation, and evaluation data and a threshold of 0.5 for the construction of the segmentation masks. Training the U-net on *NVIDIA GeForce RTX 2080* GPU yields the following results: Dice-loss of 0.095, mean accuracy of 96.2 %, mean IU of 90.9 %, and Dice-coefficient of 90.5 %.

### 2.3 Segmentation with the 3D U-net

The results for segmentation of MIP images are very encouraging. Now we consider volumetric segmentation using the 3D U-net. The resulting 3D U-net follows the same structure as shown in Figure 2.2, the only difference is the usage of 3D convolutions and 3D pooling layers.

For the 3D U-net we have to take special care about overfitting and about memory space. Therefore, for the 3D U-net we have now taken filter size 4 at the beginning and filter size 16 at the end of the contracting part. Also, the use of high dropout-rates (0.5) is necessary to ensure an efficient training process. Due to the huge size of our training samples ( $96 \times 288 \times 224$  voxels) the training of the weights takes at least half of a day. Since the number of 3D-samples is only 119, we conducted 3 training-runs with random choice of training, validation and evaluation data. Using the 3D U-net we obtained in average following results: Dice-loss of 0.219, mean accuracy of 88.3 %, mean IU of 82.4 % and Dice-coefficient of 78.2 %.

Although the 3D U-net demonstrates high precision in our application, we are not satisfied with the long training time. In addition to it, we are very limited in the choice of convolution layers and the corresponding filter sizes due to the huge size of the input data. So it is hardly possible to conduct volumetric segmentation for even larger biomedical scans without using cropping or sliding-window techniques. This is the reason why we are looking for an alternative approach.

## 3 Projection-based 2.5D U-net

One possible approach for accelerating volumetric segmentation and reducing memory requirements is to process each of the 96 slices independently through a 2D-network (compare [9]). However, this causes the loss of connection between each slice. As we have seen in Section 2.2, the 2D U-net does very well on the MIP images. So instead of processing 96 slices of size  $288 \times 224$  through a network at once, we aim for a method using some of the MIP images in connection with a learnable reconstruction algorithm.

### 3.1 Proposed 2.5D U-net architecture

A network for volumetric binary segmentation is a mapping  $\mathcal{N}: \mathbb{R}^{a \times b \times c} \rightarrow [0, 1]^{a \times b \times c}$  that maps the 3D image to the probabilities that a pixel corresponds to the desired class. In particular, the proposed 2.5D U-net takes the form

$$\mathcal{N}(x) = \mathcal{F} \circ \mathcal{M}^\sharp \circ \begin{bmatrix} \mathcal{U} \circ \mathcal{M}_1(x) \\ \vdots \\ \mathcal{U} \circ \mathcal{M}_p(x) \end{bmatrix}. \quad (3.1)$$

Here  $\mathcal{M}_i: \mathbb{R}^{a \times b \times c} \rightarrow \mathbb{R}^{b \times c}$  are MIPs,  $\mathcal{U}$  is a 2D U-net producing probabilities,  $\mathcal{M}^\sharp: (\mathbb{R}^{b \times c})^p \rightarrow \mathbb{R}^{a \times b \times c}$  is a reconstruction operator and  $\mathcal{F}: \mathbb{R}^{a \times b \times c} \rightarrow \mathbb{R}^{a \times b \times c}$  an additional filtration. We compute  $\mathcal{M}_i$  by rotating the 3D data around the vertical axis for angles  $\alpha \in \Theta$ . The 2D U-net  $\mathcal{U}$  has exactly the same structure as in

Section 2.2. The reconstruction operator  $\mathcal{M}^\sharp$  converts the output of the U-nets to 3D using linear backprojection as illustrated in Figure 3.1.




Figure 3.1: **Reconstruction operator  $\mathcal{M}^\sharp$ .** Voxel value is defined as sum over the corresponding 2D values, here illustrated for two projection angles  $\Theta = \{0^\circ, 90^\circ\}$ .

Due to the summation via  $\mathcal{M}^\sharp$ , the ideal decision boundary becomes shifted into the positive direction. We therefore subsequently apply a self-implemented layer included in  $\mathcal{F}$ , which learns the optimal back-shift during training and applies a ReLU-activation function. Because we want the output consist of sigmoid-activations, we implement another learnable layer, which shifts the neurons, inflates them to obtain more distinct decisions and applies the sigmoid-activation function. Last fine-optimization is done by a 3D-average pooling layer with pool size  $(2, 2, 2)$ . The 2D U-net and the learnable reconstruction-parameters are adjusted during the same training process.

### 3.2 Results and discussion

In our experiments, we discovered that the accuracy of the projected U-net structure mainly depends on the choice of  $\Theta$ . We decided to choose 6 equidistant angles, i.e.  $\Theta = \{k \times 30^\circ \mid k = 0, \dots, 5\}$ . We have conducted 10 training runs with random choice of training, validation and evaluation data again and obtained in average following results: Dice-loss of 0.259, mean accuracy of 86.9 %, mean IU of 79.6 % and Dice-coefficient of 74.2 %. In terms of the evaluation metrics we do not reach the



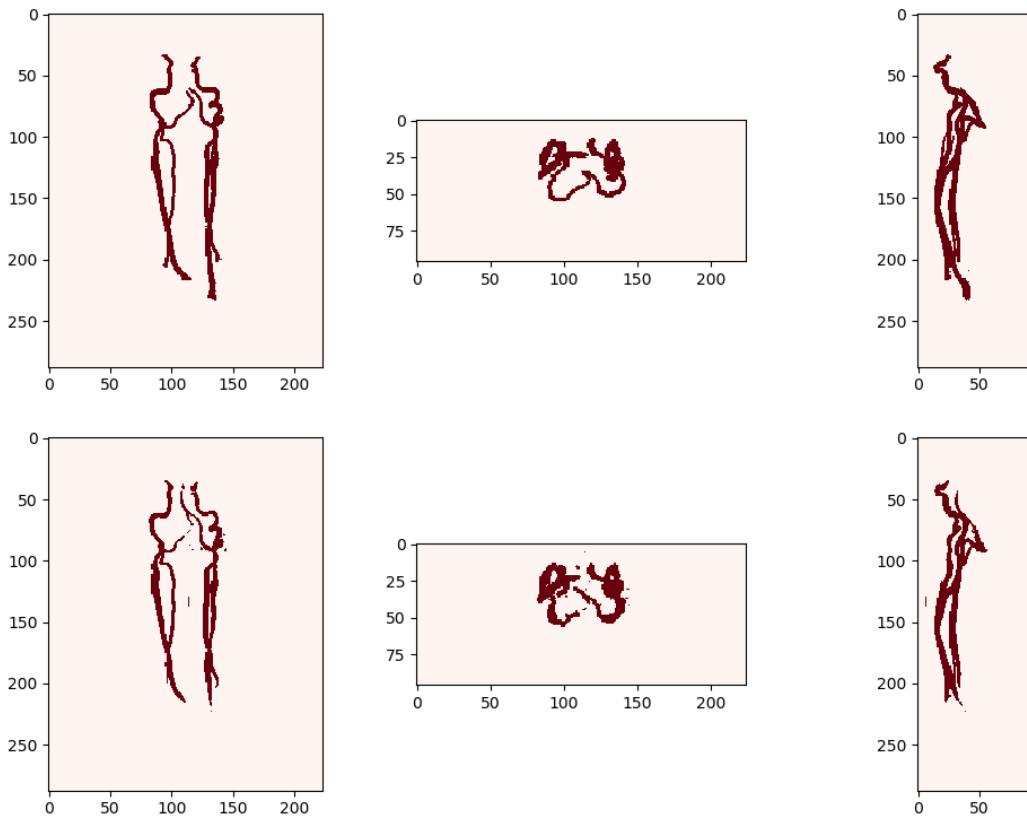


Figure 3.2: Comparison between the ground truth (first row) and the network’s segmentation (second row).

accuracy of the 3D U-net. However, we are surprised how satisfying the volumetric segmentation masks look like, see Figures 3.2 and 3.3. Note that these segmentations got produced by a network, which only uses 6 MIPs of the volumetric input data and can be trained nearly as fast as the 2D U-net. The evaluation metrics for all methods are summarized in Table 1.

Network	loss	Mean accuracy	Mean IU	Dice-coeff.
2D U-net	0.095	96.2	90.9	90.5
3D U-net	0.219	88.3	82.4	78.2
2.5D U-net	0.259	86.9	79.6	74.2

Table 1: Summarization of the numerical results.

In the current implementation, we use deterministic MIPs as input of the 2D U-nets. This causes loss of information in the case of few projection directions. In future work, we will investigate the use of more general (random) projections for  $\mathcal{M}_i$  that might contain more information in a small number of projections.




Figure 3.3: Volumetric segmentation mask generated by the 2.5D U-net.

## 4 Conclusion

In this paper we proposed a new projection-based 2.5D U-net structure for fast volumetric segmentation. The construction of volumetric segmentation masks with the help of a 3D U-net delivers very satisfying results, but the long training time and the big need of memory space are hardly sustainable. The 2.5D U-net is able to conduct reliable volumetric segmentation of very big biomedical 3D-scans and can be trained nearly as fast as a 2D-network without any concern about memory space. At the moment we do not reach the performance of the 3D U-net. However, we expect that it is possible to further increase the accuracy of 2.5D U-net, especially by modifying the MIP construction and the training process.

## Acknowledgments

C.A. and M.H. acknowledge support of the Austrian Science Fund (FWF), project P 30747-N32.

## References

- [1] J. Brownlee. Gentle introduction to the adam optimization algorithm for deep learning. <https://machinelearningmastery.com/adam-optimization-algorithm-for-deep-learning>. Last accessed: 2019-01-29.
- [2] B. Erden, N. Gamboa, and S. Wood. 3D convolutional neural network for brain tumor segmentation. Technical Report, 2018.
- [3] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016.
- [4] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [5] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015.
- [6] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In N. Navab, J. Hornegger, W. Wells, and A. Frangi, editors, *MICCAI*, pages 234–241, Cham, 2015. Springer International Publishing.
- [7] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.
- [8] P. A. Yushkevich, J. Piven, R. C. Hazlett, H. and G. Smith, S. Ho, J. C. Gee, and G. Gerig. User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *Neuroimage*, 31(3):1116–1128, 2006.
- [9] Ö. Çiçek, A. Abdulkadir, S. Lienkamp, T. Brox, and O. Ronneberger. 3D U-Net: Learning dense volumetric segmentation from sparse annotation. In *MICCAI*, 2016.