

## Gp-UNET

### Lesion detection from weak labels with a 3D regression network

Dubost, Florian; Bortsova, Gerda; Adams, Hieab H.; Ikram, M. Arfan; Niessen, Wiro J.; Vernooij, Meike W.; de Bruijne, Marleen

**DOI**

[10.1007/978-3-319-66179-7\\_25](https://doi.org/10.1007/978-3-319-66179-7_25)

**Publication date**

2017

**Document Version**

Final published version

**Published in**

Medical Image Computing and Computer Assisted Intervention – MICCAI 2017 - 20th International Conference, Proceedings

**Citation (APA)**

Dubost, F., Bortsova, G., Adams, H. H., Ikram, M. A., Niessen, W. J., Vernooij, M. W., & de Bruijne, M. (2017). Gp-UNET: Lesion detection from weak labels with a 3D regression network. In M. Descoteaux (Ed.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2017 - 20th International Conference, Proceedings* (Vol. 10435 LNCS, pp. 214-221). (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Vol. 10435 LNCS). Springer. [https://doi.org/10.1007/978-3-319-66179-7\\_25](https://doi.org/10.1007/978-3-319-66179-7_25)

**Important note**

To cite this publication, please use the final published version (if applicable).  
Please check the document version above.

**Copyright**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights.  
We will remove access to the work immediately and investigate your claim.

# GP-Unet: Lesion Detection from Weak Labels with a 3D Regression Network

Florian Dubost<sup>1,2,3(✉)</sup>, Gerda Bortsova<sup>1,2,3</sup>, Hieab Adams<sup>3,4</sup>, Arfan Ikram<sup>3,4,5</sup>,  
Wiro J. Niessen<sup>1,2,3,6</sup>, Meike Vernooij<sup>3,4</sup>, and Marleen De Bruijne<sup>1,2,3,7</sup>

<sup>1</sup> Biomedical Imaging Group Rotterdam, Erasmus MC, Rotterdam, The Netherlands  
floriandubost1@gmail.com

<sup>2</sup> Department of Medical Informatics, Erasmus MC, Rotterdam, The Netherlands

<sup>3</sup> Department of Radiology, Erasmus MC, Rotterdam, The Netherlands

<sup>4</sup> Department of Epidemiology, Erasmus MC, Rotterdam, The Netherlands

<sup>5</sup> Department of Neurology, Erasmus MC, Rotterdam, The Netherlands

<sup>6</sup> Imaging Physics, Faculty of Applied Sciences, TU Delft, Delft, The Netherlands

<sup>7</sup> Department of Computer Science,

University of Copenhagen, Copenhagen, Denmark

**Abstract.** We propose a novel convolutional neural network for lesion detection from weak labels. Only a single, global label per image - the lesion count - is needed for training. We train a regression network with a fully convolutional architecture combined with a global pooling layer to aggregate the 3D output into a scalar indicating the lesion count. When testing on unseen images, we first run the network to estimate the number of lesions. Then we remove the global pooling layer to compute localization maps of the size of the input image. We evaluate the proposed network on the detection of enlarged perivascular spaces in the basal ganglia in MRI. Our method achieves a sensitivity of 62% with on average 1.5 false positives per image. Compared with four other approaches based on intensity thresholding, saliency and class maps, our method has a 20% higher sensitivity.

## 1 Introduction

This paper addresses the problem of the detection of small structures in 3D images. We aim at developing a machine learning method requiring the least possible annotations for training. Several deep learning techniques [1–3] have recently been proposed for 3D segmentation. These methods use fully convolutional networks (FCN) [4] with a downsampling and upsampling path and can therefore detect small regions and fine details. Although efforts have been made to reduce the amount of annotations required with e.g., sparse annotations [1], these techniques still need pixel-wise annotations for training. Acquiring those is very time-consuming and often not feasible for large datasets.

In [5] the authors propose 2D localization networks requiring only global image labels for their training. They aggregate the last features maps of their network into scalars thanks to a global pooling layer and can therefore compute

a classification. Heatmaps can be obtained as a weighted sum of these feature maps. These heatmaps indicate which regions of the image contributed the most to the classification results. Because of the downsampling introduced by the pooling operations in the network, the heatmaps are several times smaller than the original input image, which makes it impossible to detect very small structures. In [6] a similar technique is applied to 3D CT lung data. This network splits into two branches, one with fully connected layers for classification, and the other with a global pooling layer for localization. The loss is computed as a weighted sum of these two terms. However, as in [5], this network is not suitable for the detection of small structures.

In this paper we propose a method to learn fine pixelwise detection from image-level labels. By combining global pooling with a fully convolutional architecture including upsampling steps as in the popular U-Net [7], we compute a heatmap of the size of the input. This heatmap reveals the presence of the targeted structures. During the training, unlike [5] or [6] where the authors use a classification network, the weights of the network are optimized to solve a regression task, where the objective is to predict the number of lesions.

We evaluate our method on the detection of enlarged perivascular spaces (EPVS) in the basal ganglia in MRI. EPVS is an emerging biomarker for cerebral small vessel disease. The detection of EPVS is a challenging task - even for a human observer - because of the very small size and variable shape. In [8–10] the authors propose different EPVS segmentation methods. However detection of individual EPVS has not been addressed much in the literature. Only [9] proposes an automatic method for this, but this method requires pixel-wise annotated scans for training.

## 2 Methods

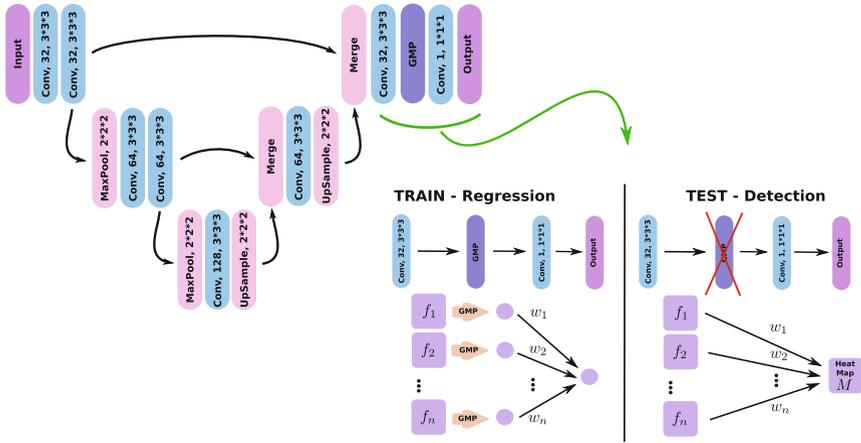
Our method takes as input a full 3D MRI brain scan, computes a smaller smooth region of interest (ROI) and inputs it to a FCN which computes a heatmap revealing the presence of the targeted brain lesions. The FCN is trained with weak labels. Its architecture changes between training and testing but the optimized weights stay the same.

### 2.1 Preprocessing

Scans are first registered to MNI space. A binary mask segmenting a region of interest (ROI) in the brain is computed with standard algorithms [11]. This mask is dilated and its borders are smoothed with a Gaussian kernel. After applying the mask, we crop the image around the ROI, trim its highest values and rescale its intensities such that the input to the network is  $S \in [0, 1]^{h \times w \times d}$ .

### 2.2 3D Regression Fully Convolutional Network

Figure 1 shows the architecture of our network. It is similar to the one of 3D U-Net [1] but is less deep and has, during training, a global pooling layer [12]



**Fig. 1. 3D Regression FCN Architecture.** Top-left: network architecture (see Sect. 2.2). Bottom-right: during training, the network is built to solve a regression problem and outputs  $\hat{y} \in \mathbb{R}$  computed according to Eq. (1). During testing, the global pooling layer is removed. Using Eq. (3), the network computes a heatmap  $M \in \mathbb{R}^{h \times w \times d}$  of the same size as the network input volume  $S$ .

before the last layer. The use of this layer is detailed in the next section. Our network has  $3 \times 3 \times 3$  convolutional layers. In the encoding path, these layers are alternated with  $2 \times 2 \times 2$  pooling layers to downsample the image and increase the size of the receptive field. In the decoding path we use  $2 \times 2 \times 2$  upsampling layers. Skip connections connect the encoding and decoding paths. We do not use padding. After each convolutional layer - except for the last one - we compute a rectified linear unit activation. The depth and number of feature maps of the network are set to fit the available GPU memory. See Fig. 1 for the architecture details.

We change the last layers of our network between training and testing. The testing step is performed with a standard fully convolutional architecture. The output is a heatmap of the size of the input. During training, we use only global image labels. To compute the loss function we need therefore to collapse the image output into a scalar. For this purpose, during training only, we introduce a global pooling layer before the last layer. In others word, we train a regression network and transform it to a segmentation network for testing (See Fig. 1). We detail this below.

**Training - Regression Network.** After the last convolutional layer, instead of stacking the voxels of each feature map and feeding them to a fully connected layer, we use a global pooling layer, which computes one pooling operation for each feature map. The resulting output of such a layer is a vector of scalars  $x \in \mathbb{R}^n$ , with  $n \in \mathbb{N}$  the number of features maps of the preceding layer.

Let  $f_i \in \mathbb{R}^{h \times w \times d}$ , for  $i \in \{1, 2, \dots, n\}$ , be the  $i$ -th feature map received by the global pooling layer. We can write the global pooling operation as  $g(f_i) = x_i$ , with  $x_i \in \mathbb{R}$ . Let  $G$  be the underlying mapping of the global pooling layer and  $f \in \mathbb{R}^{n \times h \times w \times d}$  the vector containing the  $n$  feature maps  $f_i$ . We have  $G(f) \in \mathbb{R}^n$  and in the case of global max pooling we can write  $(G(f))_i = \max_{x,y,z} f_{i,x,y,z}$ .

A convolution layer with a single output feature map  $\hat{y}$  and a filter size of  $1 \times 1 \times 1$  can be written as

$$\hat{y} = \sum_i w_i h_i + b, \quad (1)$$

with  $w_i \in \mathbb{R}$ , the weights,  $b$  the bias and  $h_i$  the  $i$ -th feature map of the preceding layer. Considering the layer following the global pooling layer, we can write Eq. (1) as

$$\hat{y} = \sum_i w_i (G(f))_i + b = w \cdot G(f) + b, \quad (2)$$

with  $\cdot$  denoting the scalar product,  $\hat{y} \in \mathbb{R}$ ,  $b \in \mathbb{R}$  and  $w \in \mathbb{R}^n$ . The output of the network  $\hat{y}$  is the estimated lesion count. There is no activation function applied after this last convolutional layer. To optimize the weights, we compute a root mean square loss  $l = \sqrt{\frac{1}{m} \sum_{t=1}^m (\hat{y}_t - y_t)^2}$ , where  $\hat{y}_t \in \mathbb{R}$  is the predicted count,  $y_t \in \mathbb{N}$  the ground truth for volume  $S_t$  and  $m$  the number of samples.

**Testing - Detection Network.** During testing we remove the global pooling layer, but we do not change the weights of the network. The  $h_i$  of Eq. (1) are the feature maps  $f_i$  defined earlier and the bias  $b^l$  can be omitted. We can thus rewrite Eq. (2) as

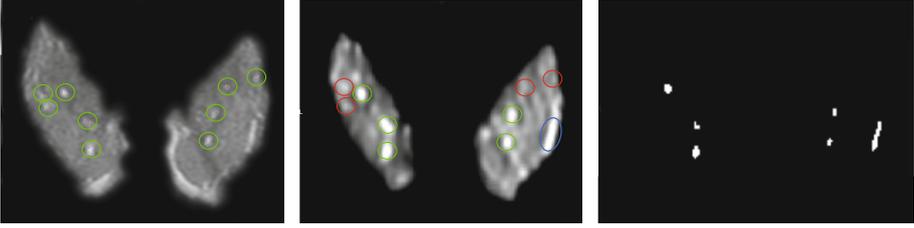
$$M = \sum_i w_i f_i, \quad (3)$$

with  $M \in \mathbb{R}^{h \times w \times d}$  the new output of the network.  $M$  is a heatmap indicating the location of the targeted lesions. Note that computation of  $M$  is mathematically equivalent to the class map computation proposed in [5]. However  $M$  has the same size as the input  $S$  and the weights  $w_i$  are optimized for a regression problem instead of classification.

After computing the heatmap, we need to select an appropriate threshold before we can compute a segmentation and retrieve the location of the targeted lesions. This efficiently be done by using the estimate of the lesion count  $\hat{y}$  provided by the neural network: we only need to keep the global pooling layer while testing. The threshold is then selected so that the number of connected components in the thresholded heatmap is equal to the lesion count.

### 3 Experiments

We evaluate our method on a testing set of 30 MRI scans and compute the sensitivity, false discovery rate and average number of false positive per image (Fig. 2).



**Fig. 2. Examples of EPVS detection in the basal ganglia.** Left: Ground truth on the intensities - the lesions are circled in green. Center: Heatmap from the neural network. True positives in green, false positives in blue and false negatives in red. Right: Segmentation of the heatmap.

**Data.** Our scans were acquired with a 1.5 T scanner. Our dataset is a subset of the Rotterdam Scan Study [13] and contains 1642 3D PD-weighted MRI scans. Note that in our dataset PD-weighted images have a signal intensity similar to T2-weighted images, the modality commonly used to detect EPVS. The voxel resolution is  $0.49 \times 0.49 \times 0.8 \text{ mm}^3$ . Each of the scans of dataset is annotated with a single, global image label: a trained observer counted the number of EPVS in the slice of the basal ganglia showing the anterior commissure. In a random subset of 30 scans, they marked the center of all EPVS visible in this same slice. These 30 scans are kept for testing set. The remaining dataset is randomly split into the following subsets: 1289 scans for training and 323 for validation.

**Experimental Settings.** The initial ROI segmentation of the BG is computed with the subcortical segmentation of FreeSurfer [11]. For registration to MNI space we use the rigid registration implemented in Elastix [14] with the mutual information as similarity measure and default settings. The Gaussian blurring kernel has a standard deviation  $\sigma = 2$ . The preprocessed image  $S$  has a size of  $168 \times 128 \times 84$  voxels. We trim its 1% highest values. We initialize the weights of the CNN by sampling from a Gaussian distribution, use Adadelta [15] for optimization and augment the training data with randomly transformed samples. A lesion is counted as detected if the center of a connected component in the thresholded heatmap is located within a radius of  $x$  pixels of the annotation. In our experiments we set  $x = 3$  which corresponds to the average radius of the largest lesions. We implemented our algorithms in Python with Keras and Theano libraries and ran the experiments on a Nvidia GeForce GTX 1070 GPU.

**Baseline Methods.** We compare our method with four conventional approaches, using the same preprocessing (Sect. 2.1). The first method (a) thresholds the input image based on its intensities: EPVS are among the most hyperintense structures in the basal ganglia. Both second and third methods compute saliency maps [16] using a neural network. Saliency maps are obtained by computing the gradient of the regression score with respect to the input image. These maps can be easily computed given any classification or regression

**Table 1. Results of our method in comparison to the baselines.** TPR stands for true positive rate (i.e. sensitivity), FPav is the number of false positives per image in average and FDR stands for false discovery rate. The comparing methods are described in Sect. 3.

Method	TPR	FPav	FDR
Intensities (a)	40.6	2.3	59.8
Saliency (b)	39.8	2.7	54.2
Saliency FCN (c)	18.7	3.3	70.2
Regression (d)	19.6	3.2	70.5
Regression FCN (e)	54.8	1.9	37.7
Intensities + Reg FCN (f)	<b>62.0</b>	<b>1.5</b>	<b>31.4</b>

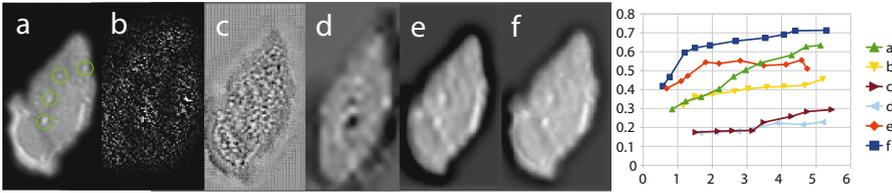
network architecture. The third method (c) uses the same network architecture as ours, the second one (b) uses a regression network without the upsampling path. Finally we compared with a method similar to [5] where we replace the original classification objective with a regression. This network is the same as ours but without the upsampling part. The thresholds are chosen based on the lesion count estimated by each network as explain in Sect. 2.

**Results.** In Table 1 we report sensitivity, false discovery rate and average false positive per image (FPav). Our method (e) detects 54.8% of the lesions with on average 1.9 false positive per image. We found that the performance could be improved further by thresholding the weighted sum of the heatmap and the original voxel intensities of the image (f). This improves the sensitivity to 62.0% with 1.5 FPav. Considering the sensitivity, our method outperforms the other baseline methods by more than 20% (Table 1). By modifying the heatmap threshold to overestimate the lesion count, we can increase the sensitivity further - 71.1% - at the price of more false positive - 4.4 in average and a false discovery rate of 60% (Fig. 3).

## 4 Discussion

Our experiments show that using a fully convolutional network and optimizing its weights for a regression problem gives the best results on our dataset (Fig. 3 and Table 1).

The methods in [5,6] produce coarse localization heatmaps on which it is impossible to locate small structures. In our early experiments we computed similar networks, optimized for classification. The resulting heatmaps highlight the whole ROI, which does not help to solve the detection problem. Using a regression network, the more precise output allows us to extract finer-scale information, but without an upsampling path, the results are still imprecise (see Fig. 3 method (d)). In our experiments, saliency maps [16] perform similarly to voxel intensity. By construction these maps are noisy (Fig. 3 method (b)).



**Fig. 3. Comparison with baselines.** Left: Heatmaps of the different methods. EPVS are circled in green on the original image (a). The saliency maps (b, c(FCN)) are very noisy. Our method (e) has finer details than a network without upsampling path (d) similar to [5]. Combining the intensities with our heatmap (f) offers the best results. Right: Free-ROC curves for each method, showing the FPav on the x-axis and the sensitivity on the y-axis.

Considering other works in automatic detection of EPVS, in [9] the authors compute an EPVS segmentation using a random forest classifier. They train this classifier with pixelwise ground truth annotations. They also use a 7 T scanner which produces high resolution images and eases small lesion detection. They report a sensitivity of 78% and a false discovery rate of 55%. This is better than our 71% sensitivity and 60% false discovery rate. However contrary to [9] which requires pixelwise annotation, our method uses only weak labels for training. In [8, 10], the authors also compute segmentations, but the evaluation is limited to a correlation with a five-category visual score.

Overall we consider that our detection results are satisfactory, considering that, during training, our method did not use any ground truth information on the location or shape of the lesions, but only their count within the whole ROI.

## 5 Conclusion

We presented a novel 3D neural network to detect small structures from global image-level labels. Our method combines a fully convolutional architecture with global pooling. During training the weights are optimized to solve a regression problem, while during testing the network computes lesion heatmaps of the size of the input. We demonstrated the potential of our method to detect enlarged perivascular spaces in brain MRI. We obtained a sensitivity of 62% and on average 1.5 false positives per scan. Our method outperforms four other approaches with an increase of more than 20% in sensitivity. We believe this method can be applied to the detection of other brain lesions.

**Acknowledgments.** This research was funded by The Netherlands Organisation for Health Research and Development (ZonMw) Project 104003005.

## References

1. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D U-Net: learning dense volumetric segmentation from sparse annotation. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016. LNCS, vol. 9901, pp. 424–432. Springer, Cham (2016). doi:[10.1007/978-3-319-46723-8\\_49](https://doi.org/10.1007/978-3-319-46723-8_49)
2. Chen, H., Dou, Q., Yu, L., Qin, J., Heng, P.A.: VoxResNet: deep voxelwise residual networks for volumetric brain segmentation. *NeuroImage* (2017)
3. Bortsova, G., van Tulder, G., Dubost, F., Peng, T., Navab, N., van der Lugt, A., Bos, D., de Bruijne, M.: Segmentation of intracranial arterial calcification with deeply supervised residual dropout networks. In: Descoteaux, M., et al. (eds.) MICCAI 2017, Part III. LNCS, vol. 10435, pp. 359–367. Springer, Cham (2017). doi:[10.1007/978-3-319-66179-7\\_41](https://doi.org/10.1007/978-3-319-66179-7_41)
4. Long, J., Shelhamer, E. and Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR (2015)
5. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A. and Torralba, A.: Learning deep features for discriminative localization. In: CVPR (2016)
6. Hwang, S., Kim, H.-E.: Self-transfer learning for weakly supervised lesion localization. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016. LNCS, vol. 9901, pp. 239–246. Springer, Cham (2016). doi:[10.1007/978-3-319-46723-8\\_28](https://doi.org/10.1007/978-3-319-46723-8_28)
7. Ronneberger, O., Fischer, P., Brox, T.: Convolutional networks for biomedical image segmentation. In: MICCAI, pp. 234–241 (2015)
8. Ramirez, J., Berezuk, C., McNeely, A.A., Scott, C.J., Gao, F., Black, S.E.: Visible Virchow-Robin spaces on magnetic resonance imaging of Alzheimer’s disease patients and normal elderly from the Sunnybrook Dementia Study. *J. Alzheimers Dis.* **43**(2), 415–424 (2015)
9. Park, S.H., Zong, X., Gao, Y., Lin, W., Shen, D.: Segmentation of perivascular spaces in 7T MR image using auto-context model with orientation-normalized features. *NeuroImage* **134**, 223–235 (2016)
10. Ballerini, L., Lovreglio, R., Hernandez, M., del C. Valdés Hernández, M., Maniega, S.M., Pellegrini, E., Wardlaw, J.M.: Application of the ordered logit model to optimising frangi filter parameters for segmentation of perivascular spaces. *Procedia Comput. Sci.* **90**, 6167 (2016)
11. Desikan, R.S., Sgonne, F., Fischl, B., Quinn, B.T., Dickerson, B.C., Blacker, D., Buckner, R.L., Dale, A.M., Maguire, R.P., Hyman, B.T., Albert, M.S.: An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage* **31**(3), 968–980 (2006)
12. Lin, M., Chen, Q., Yan, S.: Network in network. In: ICLR (2014)
13. Hofman, A., Brusselle, G.G., Murad, S.D., van Duijn, C.M., Franco, O.H., Goedegebure, A., Ikram, M.A., Klaver, C.C., Nijsten, T.E., Peeters, R.P., Stricker, B.H.C.: The Rotterdam Study: 2016 objectives and design update. *Eur. J. Epidemiol.* **30**(8), 661–708 (2015)
14. Klein, S., Staring, M., Murphy, K., Viergever, M.A., Pluim, J.P.W.: Elastix: a toolbox for intensity based medical image registration. *TMI* **29**(1), 196–205 (2010)
15. Zeiler, M.D.: ADADELTA: an adaptive learning rate method. arXiv preprint. [arxiv:1212.5701](https://arxiv.org/abs/1212.5701)
16. Simonyan, K., Vedaldi, A., Zisserman, A.: Deep inside convolutional networks: Visualising image classification models and saliency maps. In: CVPR (2014)