

## Weakly supervised object detection with 2D and 3D regression neural networks

Dubost, Florian; Adams, Hieab; Yilmaz, Pinar; Bortsova, Gerda; Tulder, Gijs van; Ikram, M. Arfan; Niessen, Wiro; Vernooij, Meike W.; Bruijne, Marleen de

**DOI**

[10.1016/j.media.2020.101767](https://doi.org/10.1016/j.media.2020.101767)

**Publication date**

2020

**Document Version**

Final published version

**Published in**

Medical Image Analysis

**Citation (APA)**

Dubost, F., Adams, H., Yilmaz, P., Bortsova, G., Tulder, G. V., Ikram, M. A., Niessen, W., Vernooij, M. W., & Bruijne, M. D. (2020). Weakly supervised object detection with 2D and 3D regression neural networks. *Medical Image Analysis, 65*, Article 101767. <https://doi.org/10.1016/j.media.2020.101767>

**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.



# Weakly supervised object detection with 2D and 3D regression neural networks

Florian Dubost<sup>a,\*</sup>, Hieab Adams<sup>b</sup>, Pinar Yilmaz<sup>b</sup>, Gerda Bortsova<sup>a</sup>, Gijs van Tulder<sup>a</sup>,  
M. Arfan Ikram<sup>c</sup>, Wiro Niessen<sup>a,d</sup>, Meike W. Vernooij<sup>b</sup>, Marleen de Bruijne<sup>a,e</sup>

<sup>a</sup> Biomedical Imaging Group Rotterdam, Departments of Radiology and Medical Informatics, Erasmus MC - University Medical Center Rotterdam, The Netherlands

<sup>b</sup> Departments of Radiology and Nuclear Medicine, and Epidemiology, Erasmus MC - University Medical Center Rotterdam, The Netherlands

<sup>c</sup> Departments of Radiology, Epidemiology and Neurology Erasmus MC - University Medical Center Rotterdam, The Netherlands

<sup>d</sup> Department of Imaging Physics, Faculty of Applied Science, TU Delft, Delft, The Netherlands

<sup>e</sup> Machine Learning Section, Department of Computer Science, University of Copenhagen, Copenhagen, Denmark

## ARTICLE INFO

### Article history:

Received 2 October 2019

Revised 12 March 2020

Accepted 22 June 2020

Available online 30 June 2020

### Keywords:

Weakly-supervised

Regression

Lesion

Detection

Weak-labels

Count

Brain

Deep learning

MRI

Enlarged perivascular spaces

Perivascular spaces

## ABSTRACT

Finding automatically multiple lesions in large images is a common problem in medical image analysis. Solving this problem can be challenging if, during optimization, the automated method cannot access information about the location of the lesions nor is given single examples of the lesions. We propose a new weakly supervised detection method using neural networks, that computes attention maps revealing the locations of brain lesions. These attention maps are computed using the last feature maps of a segmentation network optimized only with global image-level labels. The proposed method can generate attention maps at full input resolution without need for interpolation during preprocessing, which allows small lesions to appear in attention maps. For comparison, we modify state-of-the-art methods to compute attention maps for weakly supervised object detection, by using a global regression objective instead of the more conventional classification objective. This regression objective optimizes the number of occurrences of the target object in an image, e.g. the number of brain lesions in a scan, or the number of digits in an image. We study the behavior of the proposed method in MNIST-based detection datasets, and evaluate it for the challenging detection of enlarged perivascular spaces – a type of brain lesion – in a dataset of 2202 3D scans with point-wise annotations in the center of all lesions in four brain regions. In MNIST-based datasets, the proposed method outperforms the other methods. In the brain dataset, the weakly supervised detection methods come close to the human intrarater agreement in each region. The proposed method reaches the best area under the curve in two out of four regions, and has the lowest number of false positive detections in all regions, while its average sensitivity over all regions is similar to that of the other best methods. The proposed method can facilitate epidemiological and clinical studies of enlarged perivascular spaces and help advance research in the etiology of enlarged perivascular spaces and in their relationship with cerebrovascular diseases.

© 2020 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY license. (<http://creativecommons.org/licenses/by/4.0/>)

## 1. Introduction

Weakly supervised machine learning methods are designed to be optimized with limited amounts of labelled data and are very promising for a large number of medical image analysis problems. As medical expertise is scarce and annotation time expensive, unsupervised (Schlegl et al., 2017) and weakly supervised methods

(Qi et al., 2017; Bortsova et al., 2018) are most suited to extract information from large medical databases, in which labels are often either sparse or non-existent. In this article, we use attention maps for weakly supervised detection of brain lesions. Attention maps can be computed to reveal discriminative areas for the predictions of neural networks that process images such MRI, CT or X-ray. Most attention maps computation methods have originally been designed to make deep networks more explainable (Zhang et al., 2018; Oktay et al., 2018; Zhang and Zhu, 2018; Hwang and Kim, 2016). As those methods do not require annotations for the optimization of the networks but only global labels such as biomark-

\* Corresponding author.

E-mail addresses: [f.dubost@erasmusmc.nl](mailto:f.dubost@erasmusmc.nl) (F. Dubost), [marleen.debruijne@erasmusmc.nl](mailto:marleen.debruijne@erasmusmc.nl) (M.d. Bruijne).

ers or phenotypes (Wang et al., 2019), they can also be optimized using only counting objectives such as the number of lesions in a brain region, and subsequently predict the location of these lesions during test time.

We propose a novel weakly supervised detection method, using attention maps computed from the feature maps of a segmentation network architecture optimized with global labels. By using the last feature maps of such an architecture, attention maps can be computed at full input resolution, and small structures can be detected more accurately. In this article, we focus on weak supervision with regression neural networks for counting. Regression networks have widely been optimized with local labels such as voxel coordinates (Redmon et al., 2016), distance maps (Xie et al., 2018a; 2018b) or depth maps (Laina et al., 2016).

Less frequently, regression networks have been used to predict global labels, such as age (Cole et al., 2017; Wang et al., 2019), brain lesion count (Dubost et al., 2017), pedestrian count (Seguí et al., 2015), or car count (Mundhenk et al., 2016). Other researchers have also optimized neural networks to infer count. Ren and Zemel (2017) combined a recurrent network with an attention model to jointly count and segment the target objects, but need pixel-wise ground truths for the optimization. In bioimaging, methods inferring count have often been applied to cell counting in 2D images (Lempitsky and Zisserman, 2010; Walach and Wolf, 2016; Xie et al., 2018a; Tan et al., 2018; Alam and Islam, 2019). These approaches are often optimized to regress distance or density maps computed from dot annotations at the center of the target objects. Instead of regressing density maps, Paul Cohen et al. (2017) performed cell counting by regressing pixel-wise labels that represent the count of cells in the neighborhood. In our approach, pixel-wise labels are not needed for training: only the image-level count are used. Earlier, Seguí et al. (2015) have also optimized networks using image-level count labels alone for digit and pedestrian count and visualized the attention of the networks. However, they did not quantify the performance of the resulting weakly supervised detection. Xue et al. (2016) performed cell counting also using regression network optimized with patch-wise cell count, computed density maps, but did not quantify the performance on the pixel level. In this article, we optimize regression networks using image-level count labels, but use this as a means for detection.

We compare the proposed method to four state-of-the-art methods (Simonyan et al., 2014; Springenberg et al., 2015; Schlemper et al., 2018; Selvaraju et al., 2017). Other weakly supervised detection methods have been proposed relying, for example, on latent support vector machines (SVMs) (Felzenszwalb et al., 2010), a reformulation of the multiple instance learning mi-SVMs (Andrews et al., 2003), or more recently, on multiple instance learning with attention-based neural networks (Ilse et al., 2018), and on iterative learning with neural networks classifiers, where the training set is made of subsets of most reliable bounding boxes from the last iteration Sangineto et al. (2018).

We evaluate the methods using two datasets: a MNIST-based detection dataset and a dataset for the detection of enlarged perivascular spaces, a type of brain lesion that is associated with cerebral small vessel disease. On 1.5T scans, perivascular spaces become visible when enlarged. Following the neuroimaging standards proposed by Wardlaw et al. (2013), we use the consensus term perivascular space (PVS) throughout the manuscript without always referring to their enlargement. PVS is an emerging biomarker, and ongoing research attempts to better understand their etiology and relation with neurological disorders (Adams et al., 2014; Duperron et al., 2019; Gutierrez et al., 2019). Most of the research on perivascular spaces is based on quantification of PVS burden using visual scores based PVS counts (Adams et al., 2014; Potter et al., 2015). Next to overall PVS burden, the location of PVS can

have a clinical significance that varies depending on the brain region (midbrain, hippocampi, basal ganglia and centrum semiovale) and also within a brain region. For example PVS are thought to be benign when observed where perforating vessels enter the brain region (Jungreis et al., 1988), such as PVS in the lower half of the basal ganglia. Understanding more precisely how the specific locations of PVS can relate with determinants of PVS and outcomes can aid neurology research. Automatically quantifying and detecting PVS is challenging, because PVS are very small (at the limit of the scan resolution) and can easily be confused with several other types of lesions (Dubost et al., 2019b; Adams et al., 2013; Sudre et al., 2018; Brown et al., 2018). Recently, automated methods have been developed to address PVS quantification (Ballerini et al., 2018; Sudre et al., 2018; Sepehrband et al., 2019; Boespflug et al., 2018), but these methods were not evaluated in large datasets or for the detection of individual PVS. The proposed method only requires PVS visual scores for its optimization and is evaluated for the detection of individual PVS. In most of the large imaging studies, PVS are quantified using visual scores based on counts. Considering the generalizability issues of neural networks, using networks that require only PVS count for their optimization can consequently be considered to have more practical impact than networks that require annotations for their optimization.

### 1.1. State-of-the-art for attention map computation

All state-of-the-art methods investigated in this article are based on convolutional neural networks (CNNs) that compute a pseudo-probability map which indicates the locations of the target objects in the input image. In the rest of the article, we call this map the *attention map*. The methods can be divided into three categories: methods using class activation maps (CAMs), methods based on the gradient of the output of the network, and methods using perturbations of the input of the network.

**CAM methods** This category consists of variants of the class activation maps (CAMs) method proposed by Zhou et al. (2016). CAMs are computed from the deepest feature maps of the network. These feature maps are followed by a global pooling layer, and usually one or more fully connected layers to connect to the output of the network. CAMs are computed during inference as a linear combination of these last feature maps, weighted by the parameters of the fully connected layers learnt during training. If the last feature maps have a much lower resolution than the input – as is the case in deep networks with multiple pooling layers – the resulting attention maps can be very coarse. This is suboptimal when small objects need to be localized, or when contours need to be segmented precisely. To alleviate this issue, Dubost et al. (2017); Schlemper et al. (2018) proposed to include finer-scale and lower-level feature maps in the computation of the attention maps. Dubost et al. (2017) combined higher and lower level feature maps via skip connections and concatenation similarly to U-Net (Ronneberger et al., 2015), while Schlemper et al. (2018) used gated attention mechanisms, which rely on the implicit computation of internal attention maps. Selvaraju et al. (2017) proposed to generalize CAM to any network architecture, using weights computed with the derivative of the output. Unlike other CAM methods, the method by Selvaraju et al. (2017) does not require the presence of a global pooling layer in the network, and can be computed for any layer of the network.

**Gradient methods** Simonyan et al. (2014) proposed to compute attention maps using the derivative of a classification network's output with respect to the input image. These attention maps are fine-grained, but often noisy. Springenberg et al. (2015) reduced this noise by masking the values corresponding to negative entries of the top gradient (coming from the output of the network) in the ReLU activations. Gradients methods can be applied to any CNN.

**Perturbation methods** Perturbation methods compute attention maps by applying random perturbations to the input and observe the changes in the network output. These methods are model-agnostic, they can be used with any prediction model, not even necessarily restricted to neural networks. One of the simplest and most effective implementations of such methods was recently proposed by [Petsiuk et al. \(2018\)](#) with masking perturbations. The input is masked with a series of random smooth masks, before being passed to the network. Using a linear combination of these masks weighted by the updated network classification scores, the authors could compute attention maps revealing the location of the target object. This method relies on a mask sampling technique, where the masks are first sampled in a lower dimensional space, and then rescaled to the size of the full image. Earlier, [Fong and Vedaldi \(2017\)](#) proposed several other perturbation techniques including replacing a region with a constant value, injecting noise, and blurring the image. Perturbation methods are the most general as they can also be applied to other classifiers than CNN. We do not study perturbation models in this paper, because their optimization was more challenging than that of other methods, especially for the detection of small objects.

## 1.2. Contributions

The contribution of this work is fourfold. First, we propose a novel weakly-supervised detection method, named *GP-Unet*. The principle of the method is to use a segmentation architecture with skip connections to compute attention maps at full input resolution to help the detection of small objects. A preliminary version of this work was presented in ([Dubost et al., 2017](#)).


Second, the proposed method is compared to five previously published methods ([Dubost et al., 2017](#); [Schlemper et al., 2018](#); [Selvaraju et al., 2017](#); [Simonyan et al., 2014](#); [Springenberg et al., 2015](#)).

Third, we assess in MNIST-based ([LeCun et al., 1998](#)) datasets whether a classification or regression objective performs best for the weakly supervised detection.

Fourth, we evaluate the methods both in MNIST-based detection datasets and in the 3D detection of enlarged perivascular spaces. The MNIST datasets is used as a faster and more controlled experimental setting to study methodological differences between attention map computation methods, optimization objectives, and architectures. We evaluate the best methods in a real-world practical task with clinical relevance: the detection of PVS. The current work is the largest study to date to evaluate automated PVS detection in a large dataset (four regions and 2202 scans) using center locations of PVS.

## 2. Methods

We implemented seven methods for weakly supervised detection with CNNs: (a) *GP-Unet* (this article), (b) *GP-Unet no residual* ([Dubost et al., 2017](#)) the first proposed version of *GP-Unet*, (c) *Gated Attention* ([Schlemper et al., 2018](#)), (d) *Grad-CAM* ([Selvaraju et al., 2017](#)), (e) *Grad* ([Simonyan et al., 2014](#)), (f) *Guided-backpropagation* ([Springenberg et al., 2015](#)), and (g) an intensity thresholding method for brain datasets only. For all methods, the CNNs are designed to output a single scalar  $\hat{y} \in \mathbb{R}$  and are trained with mean squared error using only global labels: the number of occurrences of target objects  $y \in \mathbb{N}$ . Then for a given input image  $I$  the attention map  $M$  is computed at inference time. Below, we detail the computation of these attention maps for each method.



**Fig. 1. Principle of CAM methods for regression.** GP stands for Global Pooling.  $f_k$  correspond to the feature maps of the last convolutional layer. Disks correspond to scalar values.  $w_k$  are the weights of the fully connected layer. Left: the architecture of the network during training. Right: the architecture at inference time, where the global pooling is removed. During training, the network outputs a scalar value which is compared to the image level label to compute the loss and update the network's parameters. During testing, the global pooling layer is removed. Consequently, the network outputs an image. This image is computed as the linear combination of feature maps of the layer preceding the global pooling layer using the weights of the following fully connected layer.

## 2.1. Computation of the attention maps

### 2.1.1. CAM methods

The principle of all CAM methods is to use the feature maps – or activation maps – of the network to compute attention maps. CAM methods usually exploit the feature maps of the last convolutional layer of the network, as they are expected to be more closely related to the target prediction than feature maps of intermediate layers. [Zhou et al. \(2016\)](#) first proposed to introduce a global pooling layer after the last convolution. The global pooling layer projects each feature map  $f_k$  to a single neuron, resulting in a vector of  $N$  scalar values, where  $N$  is the number of feature maps  $f_k$  in the last layer. The global pooling layer is followed by a fully connected layer to a number of neurons corresponding to the number of classes (for classification), or to a single neuron representing the output  $\hat{y} \in \mathbb{R}$  (for regression). The network can then be trained with image-level labels using, for example, a cross-entropy or mean squared error loss function. During inference the global pooling layer can be removed, and the attention map is then computed as a linear combination of the feature maps  $f_k$  (before global pooling) using the weights of the fully connected layer  $w_k$ :

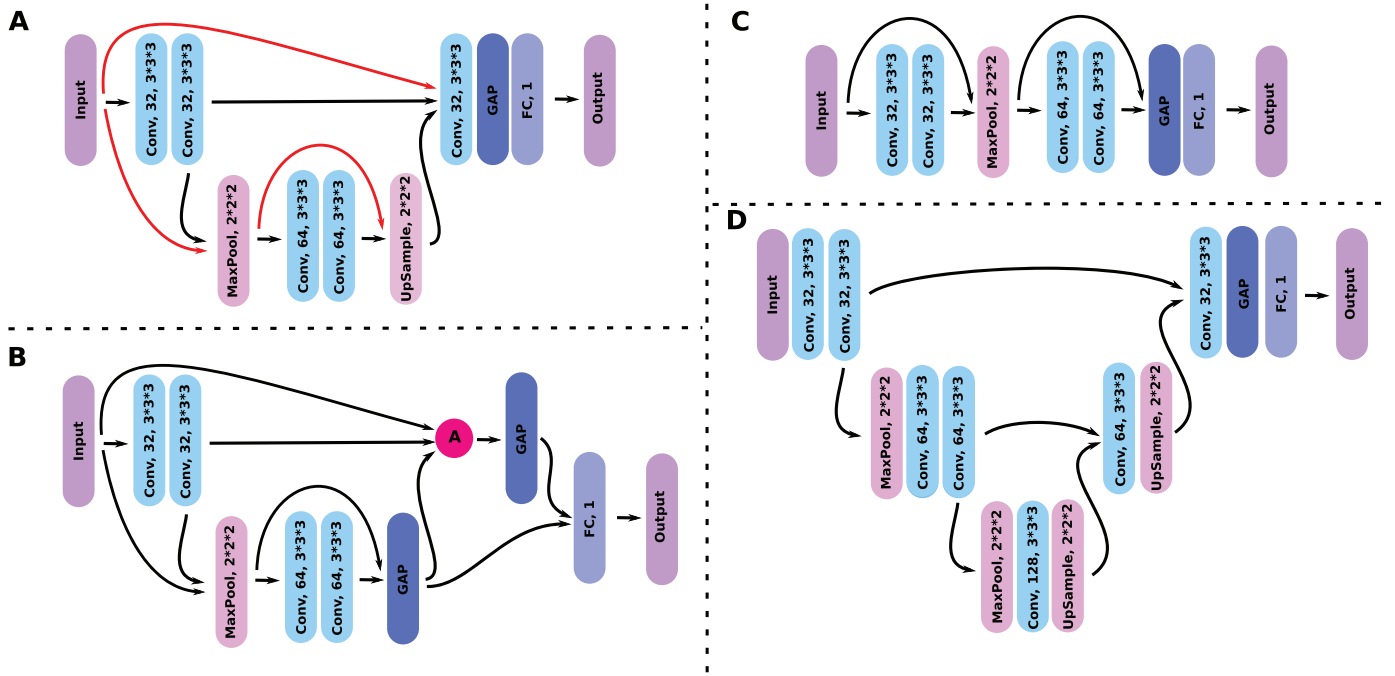
$$M_{CAM} = \sum_k^N w_k f_k. \quad (1)$$

The computation of CAM attention maps is illustrated in [Fig. 1](#).

***GP-Unet*** In the approach by [Zhou et al. \(2016\)](#) the attention map is computed from the last feature maps of the network, which are often downsampled with respect to the input image due to pooling layers in the network. To alleviate this problem, we use the same principle with the architecture of a segmentation network (U-net from [Ronneberger et al. \(2015\)](#)), i.e. with an upsampling path, where the feature maps  $f_k$  of the last convolution layer - before global pooling (GP) - have the same size as the input image  $I$  (see architectures in [Fig. 2](#) and [Section 2.2](#)). The attention maps are still computed with [Eq. 1](#).

***GP-Unet no residual***

In our earlier work, we proposed another version of *GP-Unet* ([Dubost et al., 2017](#)) based on a deeper architecture without residual connections (see architectures in [Fig. 2](#) and [Section 2.2](#)). Experiments showed that such deep architecture was not needed ([Dubost et al., 2019a](#)), and could slow the optimization. We



**Fig. 2. Architectures.** A is GP-Unet's architecture. B is Gated Attention architecture. C is the base architecture used for Grad, Guided-backpropagation, and Grad-CAM. D is GP-Unet no residual architecture. GAP stand for global average pooling layer, FC for fully connected layer, and A for attention gate. All architectures are detailed in Section 2.2. In architecture A, we showed in red the blockwise skip connections. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

refer to this approach as *GP-Unet no residual* in the rest of the paper. To detect hyperintense brain lesions in MRI data Dubost et al. (2017) also rescaled the attention map values to [0,1] and summed them pixel-wise with rescaled image intensities. This is not needed in the new version of GP-Unet above because residual connections between the input and output of two successive convolutional layers allow the network to learn this operation.

**Gated Attention** While we proposed to upsample and concatenate features maps of different scales (Dubost et al., 2017) as advised for segmentation networks by Ronneberger et al. (2015), Schlemper et al. (2018) proposed instead a more complex gated attention mechanism to combine information from different scales. This gated attention mechanism relies on attention units – also called attention gates – that compute soft attention maps and use these maps to mask irrelevant information in the feature maps. Here, global pooling is applied at every scale  $s$  and the results are directly linked to the output by a fully connected layer aggregating information across scales. Schlemper et al. (2018) proposed three aggregation strategies: concatenation, deep supervision (Lee et al., 2015), and fine-tuning by training the network for each scale separately. With the fine tuning strategy, the authors reached a slightly higher performance than concatenation and deep supervision. For the sake of simplicity, we employed the concatenation strategy in our experiments. See Fig. 2 for an illustration of the architectures of Gated Attention and of GP-Unet. The attention maps  $M_{Gated}$  of the gated attention mechanism method are computed as:

$$M_{Gated} = \sum_s \sum_k^{N_s} w_k^s f_k^s, \quad (2)$$

where  $w_k^s$  are the weights of the last fully connected layer for the neurons computed from the feature maps  $f_k^s$  at scale  $s$ .

**Grad-CAM** Finally, Grad-CAM (Selvaraju et al., 2017) is a generalization of CAM Zhou et al. (2016) to any network architecture. The computation of the attention map is similar to Eq. 1, but instead of the weights  $w_k$ , uses new weights  $\alpha_k$  in the linear combination.

The weights  $\alpha_k$  are computed with the backpropagation algorithm. With this technique the global pooling layer is not needed anymore, and attention maps can be computed from any layer in any network architecture. More precisely, each weight  $\alpha_k$  is computed as the average over all voxels of the derivative of the output  $\hat{y}$  with respect to the feature maps  $f_k$  of the target convolution layer. In our case, we use the feature maps of the last convolution layer preceding global pooling, and the weights are computed as:

$$\alpha_k = \frac{1}{Z} \sum \frac{\partial \hat{y}}{\partial f_k}, \quad (3)$$

where  $Z$  is the number of voxels in the feature map  $f_k$ . The attention map  $M_{Grad-CAM}$  is then computed as a linear combination of the feature maps weighted by the  $\alpha_k$ , and upsampled with linear interpolation to compensate the maxpooling layers:

$$M_{Grad-CAM} = \sum_k^N \alpha_k f_k. \quad (4)$$

In their original work, Selvaraju et al. (2017) proposed to compute attention maps from any layer in the network. While this approach has the advantage of generating several explanations for the network's behavior, choosing which layer should be used to compute the global attention of network becomes less obvious and objective. In our experiments, we observed that attention maps computed from the first layers of the network highlight large brain structures, and are not helpful for the detection tasks. To be more comparable to the other approaches, we used the feature maps  $f_k$  of the last convolution layer.

### 2.1.2. Gradient methods

**Grad** Simonyan et al. (2014) proposed to compute attention maps by estimating the gradient of the output with respect to the input image. Gradients are computed with the backpropagation algorithm. This method highlights pixels for which a small change would affect the prediction  $\hat{y}$  by a large amount. The attention map

$M_{Grad}$  is computed as:

$$M_{Grad} = \frac{\partial \hat{y}}{\partial I}. \quad (5)$$

**Guided-backpropagation** The attention maps obtained by Grad can highlight fine detail in the input image, but often display noise patterns. This noise mostly results from negative gradients flowing back in the rectified linear unit (ReLU) activations. In theory these negative gradients should relate to negative contributions to the network prediction, in practice they deteriorate attention maps and are believed to interact with positive gradients according to an interference phenomenon (Korbar et al., 2017). With the standard backpropagation algorithm, during the backward pass, ReLU nullifies gradients corresponding to negative entries of the bottom data (input of the ReLU coming from the input to the CNN), but not those that have a negative value in the top layer (which precedes the ReLU during the backward pass). Springenberg et al. (2015) proposed to additionally mask out the values corresponding negative entries of the top gradient in the ReLU activations. This is motivated by the deconvolution approach, which can be seen as a backward pass through the CNN where the information passes in reverse direction through the ReLU activations (Simonyan et al., 2014; Springenberg et al., 2015). Masking out these negative entries from the top layer effectively clears the noise in the attention maps.

### 2.1.3. Intensity method – for brain datasets only

PVS appear as hyperintense areas in the T2-weighted images. In some regions – especially midbrain, and to some extent basal ganglia – the image intensity can often be discriminative enough and can be used as a crude attention map. We therefore include the raw image intensity as one of the attention maps in our comparison, and, after non-maximum suppression, use the lesion count  $n$  predicted using the base architecture (see Section 2.2) to select the threshold.

## 2.2. Architectures

In total, four architectures were implemented to evaluate all six methods. These architectures are illustrated in Fig. 2. Grad, Guided-backpropagation, and Grad-CAM use the same neural networks (same architecture and weights), but differ in the computation of the attention maps during inference. The other methods require different architectures, and are trained separately. In the following section, we detail the components of each architecture in 3D.

We perform experiments on 2D CNNs for the MNIST dataset and on 3D CNNs for the brain dataset. The 3D CNNs use 3D convolutional layers with 3x3x3 filters with zero-padding, and 3D max-pooling layers of size 2x2x2. Similarly, the 2D CNNs use 2D convolutional layers with 3x3 filters with zero-padding, and 2D max-pooling layers of size 2x2. The 2D CNNs always use four times fewer feature maps than their 3D counterpart to allow faster experimentation. After the last convolution layer, each feature map is projected to a single neuron using global average pooling. These neurons are connected with a fully connected layer to a single neuron indicating the output of network  $\hat{y} \in \mathbb{R}$ . Rectified linear unit (ReLU) activations are used after each convolution. We use skip connections by concatenating the feature maps of different layers (and not by summing them).

### GP-Unet architecture (A in Fig. 2)

GP-Unet architecture is that of small segmentation network, with an encoder and a decoder part. The architecture starts with two convolutional layers with 32 filters each. The output of these two layers is concatenated with the input. Then follows a max-pooling layer and two convolutional layers with 64 filters each.

The feature maps preceding and following these two layers are concatenated. In order to combine of features at different scales, these low dimension feature maps are upsampled, concatenated with features maps preceding the maxpooling layer, and given to a convolutional layers of 32 filters. Then follows a global average pooling layer, from which a fully connected layer maps to the output. This architecture is simple (308 705 parameters for the 3D version), fast to train (less than one day on 1070 Nvidia GPU), and allows computing attention maps at the full resolution of the input image.

### GP-Unet no residual architecture (D in Fig. 2)

The architecture of GP-Unet no residual was proposed by (Dubost et al., 2017). In this work, we only changed the global pooling layer from maximum to average to make comparisons between methods more meaningful. This network is a segmentation network with a downsampling and upsampling path. The downsampling path has two convolutional layers of 32 filters, a max-pooling layer, two convolutional layers of 64 filters, a max-pooling layer, and one convolutional layer of 128 filters. The upsampling path starts with an upsampling layer, concatenates the upsampled feature maps with the features maps preceding the maxpooling layer in the downsampling path, computes a convolutional layer with 64 filters, and repeat this complete process for the last scale of feature maps, with a convolutional layer of 32 filters. After that, comes the global pooling layer, and fully connected layer to a single neuron.

The difference with architecture (A) (Dubost et al., 2017) is that the feature maps are downsampled twice instead of once, and that there are no skip connections between sets of two consecutive convolutions (blockwise skip connection in red in Fig. 2). Consequently, the last convolution layer does not have access to the input image intensities. We believe these residual connections make the design of GP-Unet more flexible than this architecture, by facilitating for instance the network to directly use the input intensities and locally adjust its predictions. This can be crucial for the correct detection of brain lesions. This architecture has twice more parameters (637 185 parameters for the 3D version) than that of GP-Unet.

**Gated Attention architecture (B in Fig. 2)** We adapted the architecture of the Gated Attention network proposed by Schlemper et al. (2018) to make it more comparable to the other approaches presented in the current work. Here, the Gated Attention architecture is the same as GP-Unet architecture (A) except for two differences: to merge the feature maps between the two different scales, instead of upsampling, concatenation and convolution, we use the attention gate as described by Schlemper et al. (2018). The other difference is that, in this architecture (B), the downsampled feature maps are also projected to single neurons with global pooling. The neurons corresponding to the two different scales are then aggregated (using concatenation) and connected to the single output neuron with a single fully connected layer. This architecture has 198 580 parameters for the 3D version.

The attention gate computes a normalized internal attention map. In their implementation, Schlemper et al. (2018) proposed a custom normalization to prevent the attention map from becoming too sparse. We did not experience such problems and opted for the standard sigmoid normalization.

Similarly to GP-Unet, Gated Attention computes attention maps at the resolution of the input image. However it combines multi-level information with a more complex process than GP-Unet.

### Base architecture (C in Fig. 2)

The network architecture used for Grad, Guided-backpropagation, and Grad-CAM is kept as similar as possible to that of GP-Unet for better comparison of methods. It starts with two convolutional layers with 32 filters each. The output

of these two layers is concatenated with the input. Then follows a maxpooling layer and two convolutional layers with 64 filters each. The output of these two layers is concatenated with the feature maps following the maxpooling layer, and is given directly to the global average pooling layer. In other words, we apply global pooling to the original image (after maxpooling) and the feature maps after the second convolution at each scale - so on 1+32+64 feature maps. This architecture has shown competitive performance on different types of problems in our experiments (eg. in brain lesions in (Dubost et al., 2019b)). With this architecture, unlike GP-Unet, Grad-CAM produces attention maps at a resolution twice smaller than that of the input image, and could miss small target objects. This architecture has 196 418 parameters for the 3D version.

### 3. Experiments

In this work, we compare our proposed method to five weakly supervised detection methods. We use the MNIST datasets (LeCun et al., 1998) to compare regression against classification for weak supervision. We compared performance of the different methods - using regression objectives - on weakly supervised lesion detection in a large brain MRI dataset.

#### 3.1. MNIST Datasets

We construct images as a grid of 7 by 5 randomly sampled MNIST digit images. Examples are shown in Figs. 4 and 5. Each digit is uniformly drawn from the set of all training/validation/testing digits, hence with a probability 0.1 to be a target digit  $d$ . To avoid class imbalance, we adapt the dataset to each target digit  $d$  by sampling 50% of images with no occurrence of  $d$ , and 50% of images with at least one occurrence of  $d$ , resulting in ten different datasets.

#### 3.2. Brain datasets


Brain MRI was performed on a 1.5-Tesla MRI scanner (GE-Healthcare, Milwaukee, WI, USA) with an eight-channel head coil to obtain 3D T2-contrast magnetic resonance scans. The full imaging protocol has been described by Ikram et al. (2015). In total, our dataset contains 2202 brain scans, each scan being acquired from a different subject.

An expert rater annotated PVS in four brain regions: in the complete midbrain and hippocampi, and in a single slice in axial view in the basal ganglia (the slice showing the anterior commissure) and the centrum semiovale (the slice 10 cm above the top of the lateral ventricle). The annotation protocol follows the guidelines by Adams et al. (2014) and Adams et al. (2013) for visual scoring of PVS, with the difference that Adams et al. (2014) only counted the number of PVS, while in the current work, all PVS have been marked with a dot in their center. Fig. 3 shows examples of PVS in the centrum semiovale.

#### 3.3. Aim of the experiments

In the MNIST datasets, the objective is to detect all occurrences of a target digit  $d$ . During optimization, the regression objective is to count the number of occurrences of  $d$ , while the classification objective is to detect the presence of at least one occurrence of  $d$ .

In the experiments on 3D brain MRI scans, the objective is to detect enlarged perivascular spaces (PVS) in the four brain regions described in Section 3.2. For these datasets we investigate only regression neural networks. These networks are optimized using the number of annotated PVS in the region of interest as the weak global label, as proposed in our earlier work Dubost et al. (2019b).



**Fig. 3. Examples of PVS in the centrum semiovale.** This is a crop of a T2-weighted image in axial view. PVS are indicated with blue arrows. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

The location of PVS are only used for the evaluation of the detection during inference.

#### 3.4. Preprocessing

**MNIST data** We scale the image intensity values in the MNIST grid images between zero and one to ease the learning process.

**Brain scans** We first apply the FreeSurfer multi-atlas segmentation algorithm (Desikan et al., 2006) to locate and mask the mid-brain, hippocampi, basal ganglia and centrum semiovale in each scan. For each region, we then extract a fixed volume centered on the center of mass of the region. For midbrain (88x88x11 voxels), hippocampi (168x128x84 voxels) and basal ganglia (168x128x84 voxels) these cropped volumes contain the full region. The centrum semiovale is too large to fit in the memory of our GPU (graphics processing unit), so for this region we only extract the slices surrounding the slice that was scored by the expert rater (250x290x14 voxels). Consequently, we apply a smooth region mask to nullify values corresponding to other brain regions. Finally, we scale the intensity values between zero and one to ease the learning process. The preprocessing and extraction of brain regions is presented in more details in previous work (Dubost et al., 2019b).

#### 3.5. Training of the networks

All regression networks are optimized with Adadelta (Zeiler, 2012) to minimize the mean squared error between their prediction  $\hat{y} \in \mathbb{R}$  and the ground truth count  $y \in \mathbb{N}$ . The classification networks in our MNIST experiments were optimized with Adadelta and the binary cross-entropy loss function.

Weights of the convolution filters and fully connected layers are initialized from a Gaussian distribution with zero mean and unit variance, and biases are initialized to zero.

A validation set is used to prevent over-fitting. The optimization is stopped at least 100 epochs after the validation loss stopped decreasing. We select the model with the lowest validation loss. For the MNIST datasets, the models are trained on a set of 500 images (400 for training and 100 for validation). For the brain datasets, the models are trained on a set of 1202 scans (1000 for training and 202 for validation). During training, we use on-the-fly data augmentation with a random combination of random translations of up to 2 pixels in all directions, random rotations up to 0.2 radians in all directions, and random flipping in all directions. For the MNIST datasets, the batch size was set to 64. For

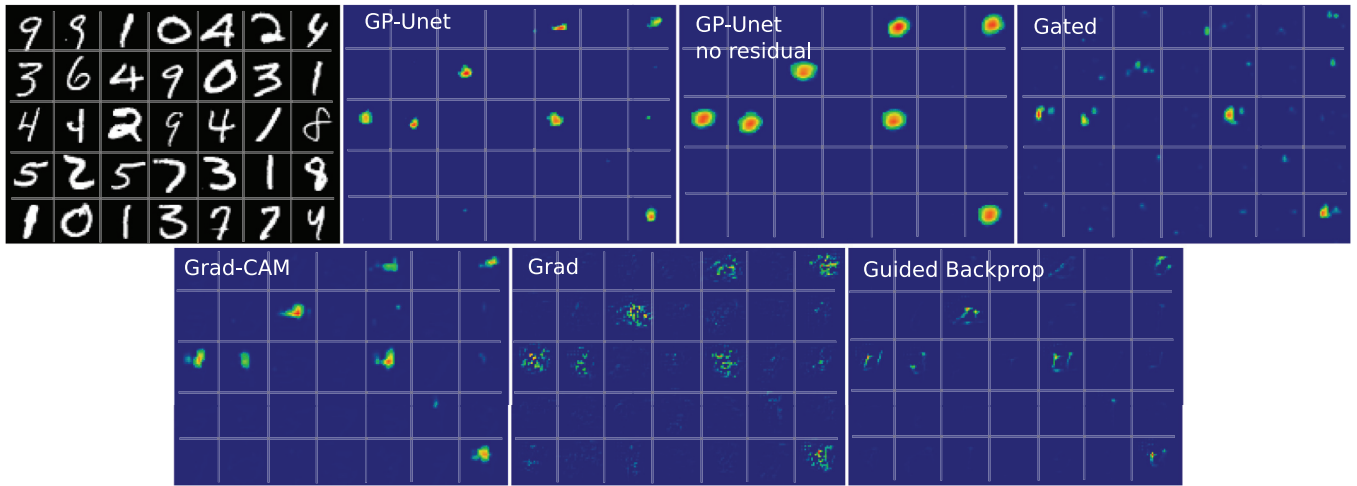


Fig. 4. Examples of attention maps of the different weakly supervised detections methods for the detection of digit 4. Top-left: MNIST image. All methods with optimized with regression objectives.




Fig. 5. Examples of attention maps of GP-Unet for the detection of digit 4 and optimized with classification and regression objectives. Left: MNIST image, middle: attention map generated from a classification network, right: attention map generated from a regression network. The first row displays an image without digit 4. The second row displays an image with seven occurrences of the digit 4. For the classification method, in the first row we notice more false positives than for the regression method. On the second row, the two digits 4 at the top are less highlighted than the other digits 4 in the image. It is not the case for the regression attention map. This observation supports the hypothesis that attention maps computed from classification objectives tend to focus more on the most obvious occurrence of the target object, instead of equally focusing on all occurrences. On the right, we show the difference between the attention maps for regression and classification.

the brain datasets, because of GPU memory constraints, the networks are trained per sample: each mini-batch contains a single 3D image. As the convergence can be slow in some datasets, we first trained the networks on the smallest and easiest region (mid-brain), and fine-tune the parameters for the other regions, similarly to Dubost et al. (2019b).

We implemented our algorithms in Python in Keras (Chollet et al., 2015) with TensorFlow as backend, and ran the experiments on a Nvidia GeForce GTX 1070 GPU and Nvidia Tesla K40<sup>1</sup>. The average training time was one day.

### 3.6. Negative values in attention maps

Attention maps can have negative values, which meaning can differ for CAM methods and gradient methods. For CAM methods, negative values could highlight objects in the image which presence is negatively associated with the target objects. For gradient methods, they correspond to areas where increasing the intensity would decrease the predicted count (or where decreasing the in-

tensity would increase the predicted count, these are the same areas).

For image understanding, keeping negative values in attention maps seems most appropriate as the purpose is to discover which parts of the image contributed either negatively or positively to the prediction, and how a change in their intensity could affect the prediction. For detection, the purpose is to find to find all occurrences of the target object in the image and ignore other objects. In the literature, two approaches have been proposed to handle negative values for object detection: either setting them to zero, or taking the absolute value. CAM methods (Zhou et al., 2016; Selvaraju et al., 2017) nullify negative values of the attention maps to mimic the behavior of ReLU activations. Gradient methods (Simonyan et al., 2014; Springenberg et al., 2015) focus on the magnitude of the derivative and thus compute the absolute value.

In our case, we aim to solve a detection problem in datasets where the target objects are among the highest intensity values in the image. For gradient methods, this implies that negatives values in the attention maps do not indicate the location of the target object in our case. We can therefore ignore negative values, and decided to nullify them. For CAM methods, we follow the recommendation of the literature, and also nullify negative values in attention maps. Consequently, we nullified negative values for all

<sup>1</sup> We used computing resources provided by SurfSara at the Dutch Cartesius cluster.



methods. Nullifying negative values actually only impacts the visualization of the attention maps, and not the detection metrics, as we select only candidates with highest values in the attention maps (Section 3.7). On the contrary taking the absolute value could increase the number of detections and would impact our detection metrics.

### 3.7. Performance evaluation

The output of all weakly-supervised detection methods presented in Section 2 are attention maps. We still need to obtain the coordinates of the detections, and evaluate the matching with the ground truth.

After setting negative values to zero (Section 3.6), we apply non-maximum suppression on the attention maps using a 2D (MNIST, centrum semiovale and basal ganglia) or 3D (hippocampi and midbrain) maximum filter of size 6 voxels (which corresponds to 3 mm in axial plane, the maximum size for PVS as defined by Adams et al. (2013) – we used the same value for the MNIST datasets) with 8 neighborhood in 2D or 26 neighborhood in 3D. This results in a set of candidates that we order according to their value in the attention map. The candidates with highest values are considered the most likely to be the target object.

For the basal ganglia and the centrum semiovale, our dataset does not contain full 3D annotations, but only provides annotations for a single 2D slice per scan (see Section 3.2). As annotations were only available in a single slice, we evaluated the attention maps only in the annotated slice, although we can compute attention maps for the complete volume of these regions. For our evaluation we extract the corresponding 2D slice from the attention map prior to post-processing and compute the metrics only for this slice. In case no lesion was annotated, we selected the middle slice of the attention map as a reasonable approximation of the rated slice.

As we aim to solve a detection problem, we need to quantify the matching between two sets of dots: the annotators dots, and the algorithms' predictions. We used the Hungarian algorithm (Kuhn, 1955) to create an optimal one-to-one match between each detected lesion or digit to the closest annotation in the ground truth. For the brain dataset, we counted a positive detection if a detection was within at most 6 voxels from the corresponding point in the ground truth. This corresponds to the maximum diameter of PVS in the axial view, as defined in Adams et al. (2013). For the MNIST datasets, we counted a positive detection if a detection fell inside the 28\*28 pixels wide original MNIST image of the target digit.

As the algorithms output candidates with confidence scores, we can compute free-response receiver operating characteristic (FROC) curves (Bandos et al., 2009) that show the trade-off between high sensitivity and the number of false positives, in our case more precisely the average number of false positives per scan (FPavg). To draw these curves, we varied the number of selected candidates. For each network in our experiments, we report the area under the FROC curve (FAUC) computed from 0 to 5 FPavg for MNIST and from 0 to 15 FPavg for brain lesion detection. We also show the standard deviation of the FAUC, computed by bootstrapping the test set.

In addition to the attention maps, the regression networks also predict the number of target objects in the image. For the detection of brain lesions, we use this predicted count rounded to an integer  $n$  to select the top- $n$  candidates with highest scores, and compute the corresponding sensitivity and FPavg, and the average number of false negative per scan (FNavg). For statistical significance of difference of FAUCs, we performed a bootstrap hypothesis testing and consider statistical significance for p-value lower than 0.05. For FPavg, FNavg and Sensitivity we performed Wilcoxon tests using p-value lower than 0.05.

### 3.8. Intra-rater variability of the lesion annotations

Intra-rater variability has been measured in each region using a separate set of 40 MRI scans acquired and annotated with the same protocol. The rater annotated PVS twice in each scan with two weeks of interval, and in a different random order.

To compute the sensitivity and FPavg for the Intra-rater variability, one of the two series of annotations has to be set as reference to define true positives, positives and false positives. We successively set the first and second series of annotations as reference, leading to two different results. All results for all regions are displayed next to the FROC curves in Fig. 7.

## 4. Results

### 4.1. Regression vs classification objectives - MNIST datasets

The methods were evaluated on left-out test sets of 500 images, balanced as described in Section 3.1. Fig. 6 compares the FAUC of regression and classification networks, for all MNIST digits, and for all methods. Additional results such as FROC curves, sensitivity, FPavg and FNavg are given in Appendix A and Appendix B. Overall, regression methods reach a higher detection performance than classification methods. For all digits, regression GP-Unet no residual reaches the best performance. The second best method for all digits is regression GP-Unet. Both GP-Unet regression methods are consistently better than any other method for all digits. Regression Grad-CAM comes third, and regression Guided-backpropagation fourth. Grad and Gated Attention come last. The ordering of best classification methods is different than that of the best (regression) methods: Guided-backpropagation comes first, Grad-CAM second and GP-Unet no residual third.


Fig. 4 shows an example of the attention maps obtained for all weakly supervised methods optimized with regression objectives. As expected, Grad produces noisy attention maps with many high values, for both classification and regression objectives, and Guided-backpropagation corrects these mistakes. Gradient methods seems to highlight multiple discriminating features of the digit 4 (e.g. its top branches), while CAM methods highlight a single larger, less detailed region. This may suggest that gradients methods may be more suited to weakly supervised segmentation, although judging from the figure, none of the methods seems capable of correctly segmenting digits.

Fig. 5 compares attention maps of GP-Unet optimized with regression and classification. We noticed two interesting differences. First, when the target digit is present on the image, the regression attention map highlights each occurrence of the target digits with a similar intensity, while the classification attention map highlights more strongly the most obvious occurrences of the target digit. Second, when the target digit is not present in the image, contrary to the regression attention map, the classification attention map may highlight many false positives, possibly resulting in a significant drop in the detection performance.

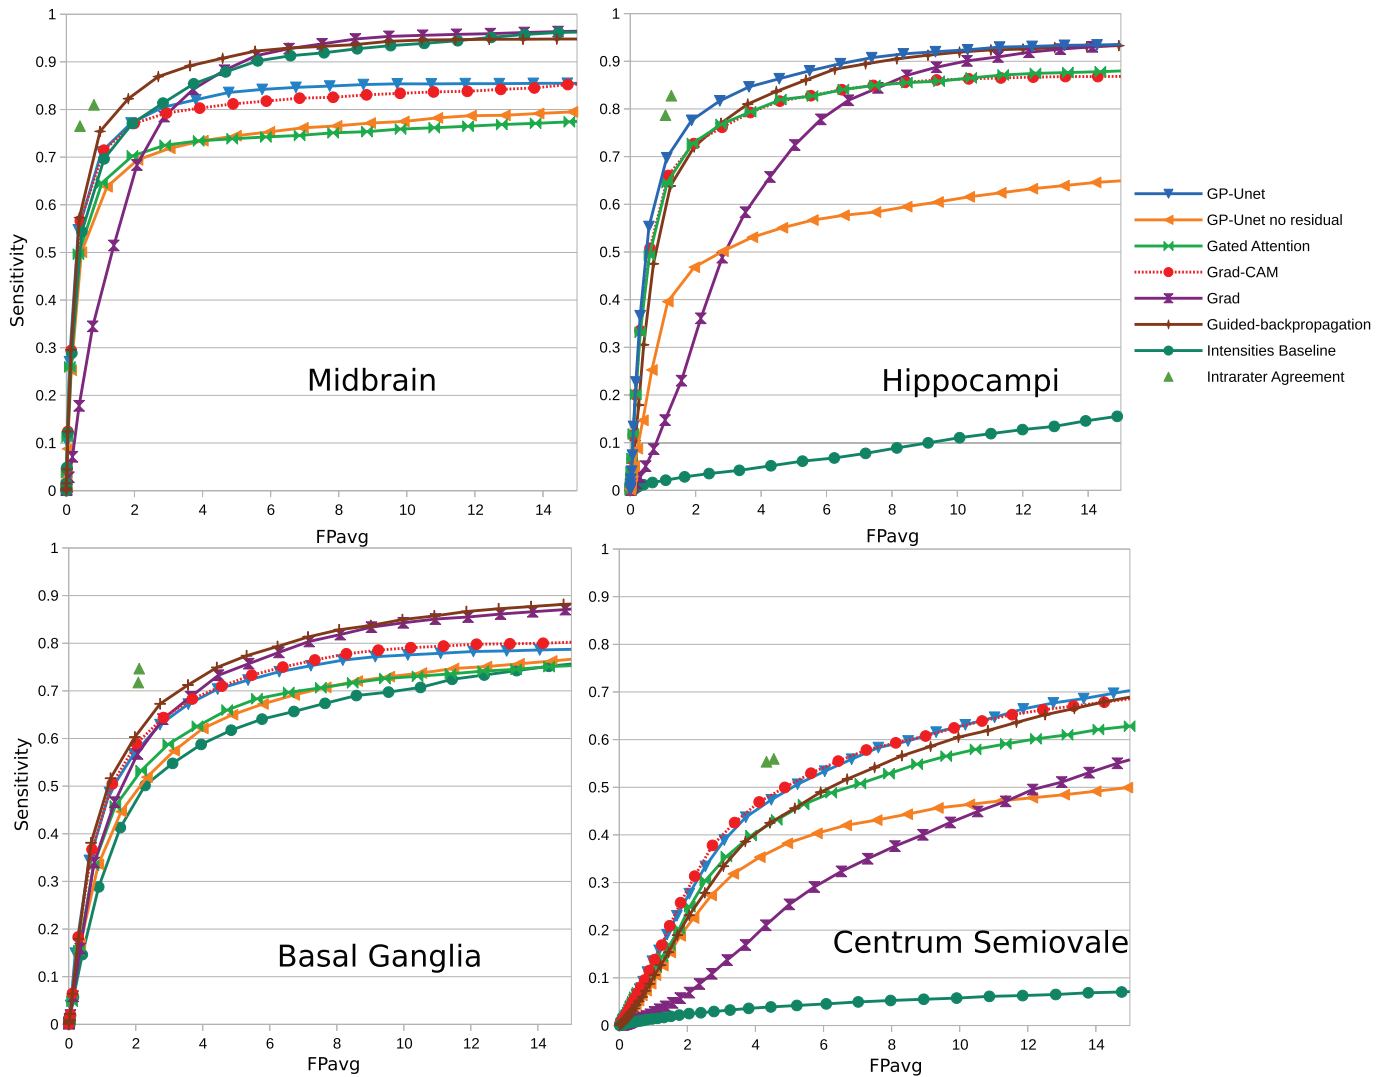
*Regression Guided-backpropagation vs Grad.* Regression Guided-backpropagation detects of all digits more accurately than regression Grad. The same comparison holds for classification Guided-backpropagation versus classification Grad. However Regression Grad sometimes performs as well (digits 4, 6, 7) or better (digits 0, 9) than Classification Guided-backpropagation, which underlines the added-value of optimizing weakly supervised detection methods with regression objectives instead of classification objectives.

### 4.2. Variations of the architecture of GP-Unet - MNIST datasets


In this section we studied the influence of the skip connections between sets of two consecutive convolutions (blockwise skip con-



**Fig. 6.** FAUCs (Section 3.7) on the MNIST dataset for all methods. Each subplot corresponds to the detection of a different digit. Results for regression networks are displayed in light blue, and results for classification networks are displayed in indigo. FAUCs are displayed with standard deviations computed by bootstrapping the test set. A is GP-Unet, B GP-Unet no residual, C Gated Attention, D Grad-CAM, E Grad and F Guided-backpropagation. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)



**Fig. 7.** FROC curves of enlarged perivascular spaces detection in the brain MRI in four different regions. The average number of false positives per scan is displayed on the x-axis, and the sensitivity on the y-axis. Axes have been rescaled for better visibility. The green triangles indicate intra-rater agreement (on a smaller set) as described in Section 3.8. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)



**Fig. 8. Attention maps in the midbrain.** The top left image shows the slice of an example image of the midbrain after preprocessing, with PVS indicated with red circles. The other images correspond to attention maps computed for that same slice. Red values correspond to high values in the attention maps. The intensity baseline method in the bottom right corner is actually the same as the image in the upper left corner but with a different color map. Values in attention maps are not bounded, and the maximum varies between images and methods. For the visualization, we chose the scaling of attention maps to best show the range of values in each image. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

nections, in red in Fig. 2) in GP-UNET's architecture and the influence of the type of global pooling in GP-UNET's architecture on the detection performance. Removing the blockwise skip connections did not make the detection worse for most digits (except digit 1 and 7 where having the blockwise skip connections helped). Using global max pooling instead of global average pooling led to worse detection performance for all digits. For all digits the optimization was better with the proposed architecture. Removing skip connections or using global max pooling made the optimization take longer to converge, made loss curves not as smooth and made the loss converged to a higher value. The corresponding FROC curves, FAUC barplot, and FAUC, FPavg, FNavg and Sensitivity Tables are given in Appendix C.

#### 4.3. Detection of brain lesions

In the brain dataset, we compare the performance of the weakly supervised methods for the detection enlarged perivascular spaces (PVS) by evaluating them on the left-out test set of 1000 scans, and in four brain regions: midbrain, hippocampi, basal ganglia, and centrum semiovale.

Figs. 8 – 11 show attention maps for all methods in the four regions. Fig. 7 shows FROC curves for all methods in the brain datasets. Table 1 shows the corresponding FAUCs. Table 2 and Table 3 show the sensitivity and FPavg measured at the operating point chosen for each method as described in Section 3.7. Table 4 shows the average number of false negatives.

Judging from Tables 1, 2, 3 and 4, the methods achieving the best results are GP-UNET, Grad-CAM and Guid-backpropagation. Unlike the results on MNIST datasets, there is no method consistently better than others for all regions. In the midbrain and basal ganglia, Guided-backpropagation reaches the best results of all methods, and in all three metrics, with the exception of FPavg in the basal ganglia. In the hippocampi, GP-UNET reaches the best results of all methods, and in all four metrics. In the centrum

semiovale, GP-UNET and Grad-CAM achieve the best results, and have a similar performance. Intensity thresholding reaches a competitive performance in the midbrain and basal ganglia, but completely fails in the hippocampi and centrum semiovale because it highlights many false positives, corresponding to other hyperintense structures. Surrounding cerebrospinal fluid, white matter hyperintensities, and sulci are examples of these structures.

In Fig. 7, the sensitivity and FPavg between two series of annotations of the same scans from the same rater (green triangle) gives an idea of the difficulty of detecting PVS in each region. In the midbrain and hippocampi, PVS are relatively easy to identify, as they are the only hyperintense lesions visible on T2 images. On the contrary, the detection of PVS in the basal ganglia and centrum semiovale is much more challenging, because in those regions other hyperintense structures that look similar to enlarged perivascular spaces. In all regions, the performance of the automated methods come close to the intra-rater agreement. This intrarater agreement was however computed on a substantially smaller set – 40 vs 1000 scans – and shorter annotation period – 1 week vs several months. Interestingly, several methods highlight the same false positives. After visual checking by experts, many of these false positives appear to be PVS annotated by the rater. In the set of 40 scans used the intrarater measures, 68 percent of false positive detections of GP-UNET in the centrum semiovale were PVS. More precisely, 39 percent of false positives were enlarged PVS and 29 percent were slightly enlarged PVS.

## 5. Discussion

Overall, results showed that weakly supervised methods can detect PVS almost as well as expert raters. The performance of the best detection methods was close to the intrarater agreement. The interrater agreement is also probably lower than this intrarater agreement. Finally, further visual inspection also revealed that many of the false positives correspond to PVS that were not




Fig. 9. Attention maps in the hippocampi.

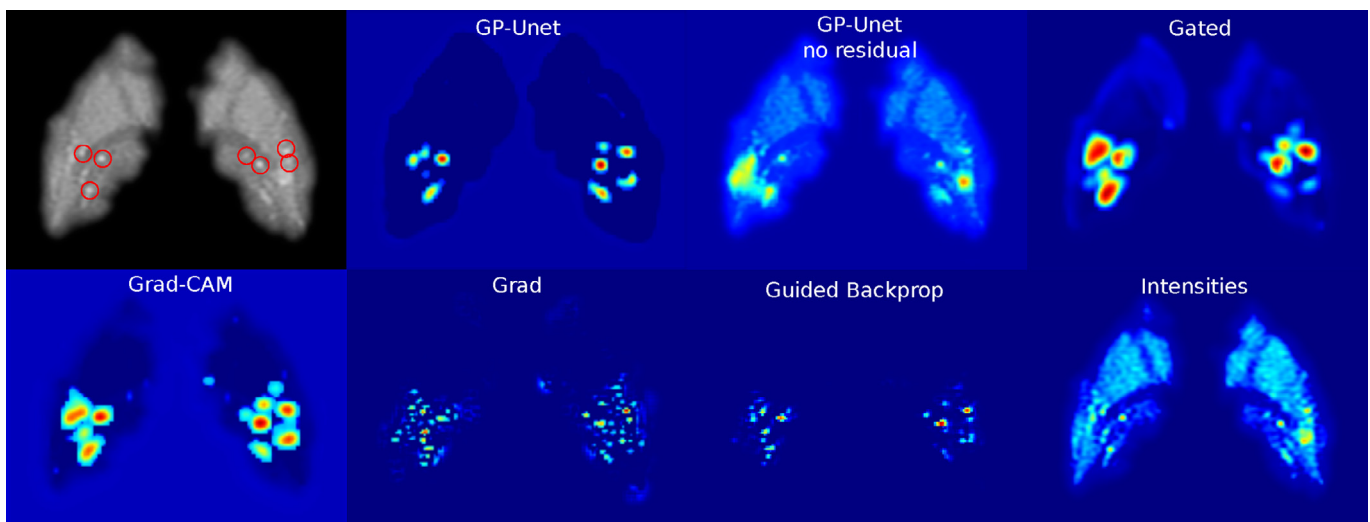


Fig. 10. Attention maps in the basal ganglia.

annotated by the human rater. We especially noticed that annotating all PVS was difficult for the expert rater in scans with many PVS.

We compared six weakly supervised detection methods in two datasets. We showed that the proposed method could be used with either 2D or 3D networks. For all methods, 2D networks in the MNIST datasets converged substantially faster (hours) than the 3D networks in the brain dataset (days). In MNIST datasets for regression, GP-UNET no residual (Dubost et al., 2017) and GP-UNET (this article) perform significantly better than all other methods, probably because they can combine the information of different scales more effectively than other methods. For GP-UNET no residual, part of this performance difference can also be explained by the larger number of parameters and larger receptive field (Section 2.2). On the contrary, for GP-UNET, the number of parameters is comparable to that of the other methods. In the brain dataset, the best methods are Guided-backpropagation (Springenberg et al., 2015) with 74.1 average FAUC over regions, GP-UNET with 72.0 average FAUC, and Grad-CAM Selvaraju et al. (2017) with 70.5 average FAUC. As GP-UNET performs either similarly to or better than Grad-CAM depending on the region, given a new weakly supervised detection

task, we would consequently recommend Guided-backpropagation and GP-UNET.

Grad-CAM and GP-UNET reach similar FAUCs (Table 1) in the basal ganglia and centrum semiovale. However, GP-UNET outperforms Grad-CAM in the midbrain and by a large margin in the hippocampi. In these two regions, at the operating point Grad-CAM suffers from more false positives than GP-UNET, while having a similar or worse sensitivity (Table 3 and 2). The attention maps of the hippocampi (Fig. C.16) – and to some extent those of the midbrain (Fig. 8) – show that GP-UNET is less distracted by the surrounding cerebrospinal fluid than Grad-CAM – or the methods emphasizing intensities (GP-UNET no residual, Intensities). The attention maps of Grad-CAM and GP-UNET share most of the false positive detections. Most of these false positives are PVS that were not annotated by the rater. Overall, the attention maps of GP-UNET are also sharper than the ones of Grad-CAM, probably because GP-UNET can compute attention maps at a higher resolution: the resolution of the input image.

The motivation of Gated Attention (Schlemper et al., 2018) is similar to that of GP-UNET: combining multiscale information in the computation of attention maps. In the MNIST datasets, while Gated Attention and GP-UNET reach a similar detection performance when

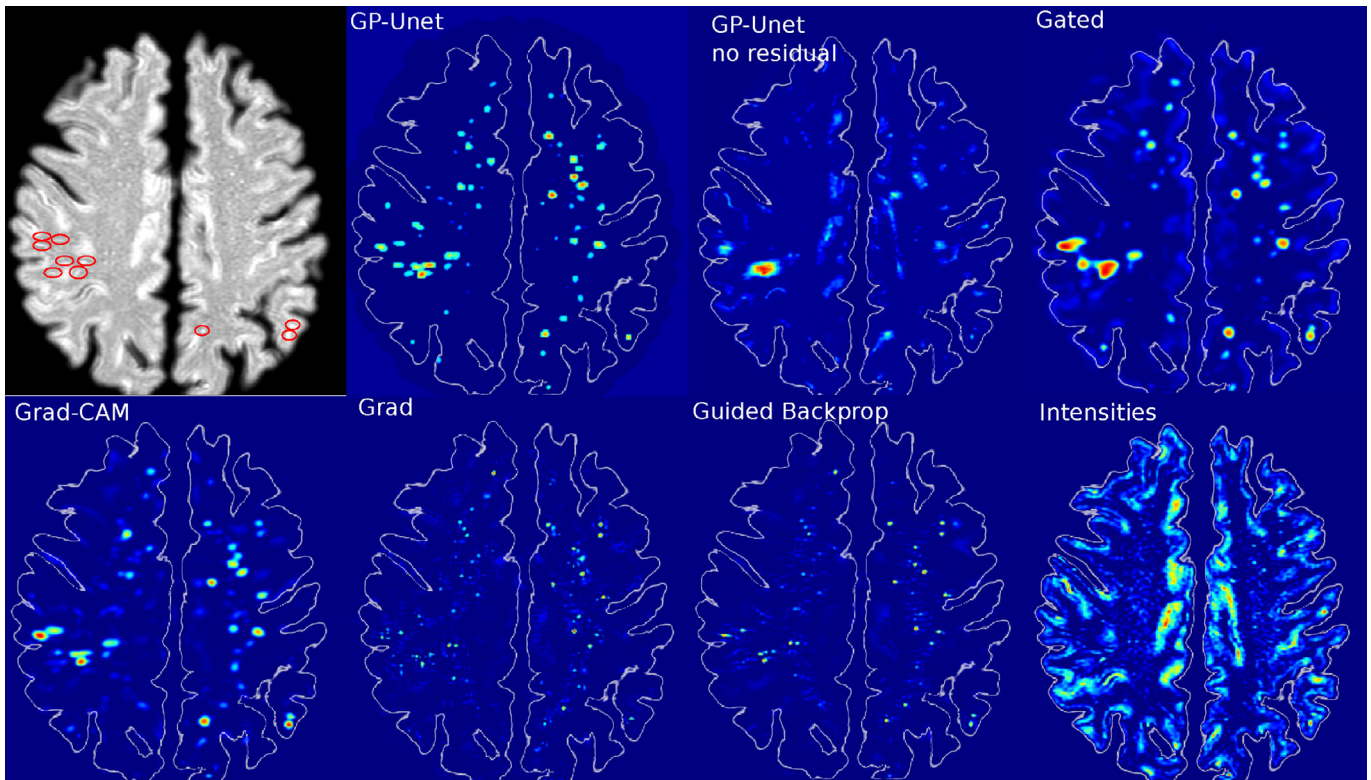


Fig. 11. Attentions maps in the centrum semiovale. Contours of the brain have been delineated in white for better visualization.

optimized with classification objectives, contrary to GP-UNET, Gated Attention rarely benefits from the regression objective. More generally, Gated Attention seems to benefit less often from the regression objective than the other methods. These results suggest that gate mechanisms may harm the detection performance for networks optimized with regression objectives, and that a simple concatenation of feature maps should be preferred. In the brain datasets, Gated Attention works better than the intensity baseline, Grad (Simonyan et al., 2014), and GP-UNET no residual, but performs significantly worse than Grad-CAM, Guided-backpropagation, and GP-UNET. One should also keep in mind that Gated Attention was originally proposed for deeper networks. In case of shallow networks, this method may not reach its full potential, as it benefits only from few (two on our case) different feature scales.

We mentioned above that the attention maps of GP-UNET are sharper than those of Grad-CAM. In Appendix C, we investigate the influence of the architecture and compare attention maps of GP-UNET, GP-UNET without blockwise skip connections (GP-UNET No Skip) and GP-UNET with global max pooling instead of global average pooling (GP-UNET Max Pool). Removing the skip connections does not seem to make the attention less compact. Using global max pooling does make the attention maps more compact but increases the number of false negatives. GP-UNET may have more compact attention maps than Grad-CAM on the basic architecture thanks to the upsampling path in GP-UNET. To compute the attention at full input resolution with Grad-CAM, the attention maps need to be interpolated, resulting in less compact attention maps. GP-UNET may have more compact attention maps than Gated Attention because concatenating feature maps might be more efficient (maybe easier to optimize) in combining multiscale features than using the gated attention.

Due to the special properties of the PVS detection problem in the brain datasets, intensity thresholding provides a simple approach to solving the same problem. Although intensity threshold-

ing yields the worst results in hippocampi, basal ganglia, and centrum semiovale, it achieves the second best FAUC in the midbrain. This high performance results from the effective region masking specific to the midbrain: because PVS are almost always in the center of this region, we can erode the border of the region mask, and eliminate the hyperintense cerebrospinal fluid surrounding the midbrain. As there are no other visible lesions in the midbrain, all remaining hyperintensities correspond to PVS.

In the datasets where the intensity method achieved good or reasonable results (midbrain and basal ganglia), Guided-backpropagation performed best. In the datasets where the intensity method failed (hippocampi and centrum semiovale), GP-UNET reached the best performance (similar to that of Grad-CAM in the centrum semiovale). More generally, gradients methods seem to work best when the target objects are also the most salient objects, while CAM methods work best when saliency alone is not discriminative enough. This observation can also be extended to the MNIST datasets, where saliency alone is not sufficient, and regression CAM methods (Gated Attention excluded) outperform regression gradient methods.

Recently Adebayo et al. (2018) showed that, for Guided-backpropagation, classification networks trained with random labels obtained similar attention maps as networks trained with the correct labels, hinting that attention maps method may focus more on salient objects in the image than the target object. In these experiments, attention maps computed with Grad and Grad-CAM obtained better results. Adebayo et al. warn of the evaluation of attention maps by only visual appeal, and advocate more rigorous forms of evaluation. This fits exactly with the purpose of the current article, in which we aimed to quantify the detection performance of attention maps in large real world datasets.

For the evaluation of the detection of PVS, images were annotated by a single rater. With the same resources, we could also have had multiple raters annotating fewer scans and use their con-

**Table 1**

**FAUCs for the detection of brain lesions.** To compute the these FAUCs, we integrate the FROC (Fig. 7) between 0 and 15 (Section 3.7). The best performance in each region is indicated in bold.

	GP-Unet (this paper)	GP-Unet no residual Dubost et al. (2017)	Gated Attention Schlemper et al. (2018)	Grad-CAM Selvaraju et al. (2017)	Grad Simonyan et al. (2014)	Guided-backprop Springenberg et al. (2015)	Intensities Section 4.3
Midbrain	81.5 (80.1 - 82.8)	73.4 (72.0 - 74.8)	72.7 (71.1 - 74.4)	79.8 (78.5 - 81.1)	84.5 (83.5 - 85.4)	<b>89.2 (88.3 - 90.2)</b>	87.1 (86.1 - 88.1)
Hippocampi	<b>85.8 (84.8 - 86.7)</b>	55.1 (53.5 - 56.7)	80.2 (79.1 - 81.3)	80.1 (78.9 - 81.3)	71.5 (70.4 - 72.6)	83.3 (82.2 - 84.3)	8.3 (7.5 - 9.0)
Basal Ganglia	69.6 (68.1 - 71.2)	64.4 (63.0 - 65.9)	64.8 (63.4 - 66.4)	70.6 (69.3 - 72.0)	<b>73.5 (72.2 - 74.9)</b>	<b>75.6 (74.3 - 76.8)</b>	61.7 (59.9 - 63.5)
Centrum Semiovale	<b>51.3 (50.1 - 52.6)</b>	37.9 (36.8 - 39.2)	46.2 (45.0 - 47.5)	<b>51.5 (50.2 - 52.7)</b>	31.9 (30.7 - 33.2)	48.1 (46.9 - 49.3)	4.7 (4.2 - 5.2)
Average	72.0 +/- 13.3	57.7 +/- 13.1	66.0 +/- 12.7	70.5 +/- 11.6	65.4 +/- 19.9	74.1 +/- 15.7	40.5 +/- 35.2

**Table 2**

**Sensitivity in the brain datasets.** Best performance are indicated in bold.

	GP-Unet (this paper)	GP-Unet no residual Dubost et al. (2017)	Gated Attention Schlemper et al. (2018)	Grad-CAM Selvaraju et al. (2017)	Grad Simonyan et al. (2014)	Guided-backprop Springenberg et al. (2015)	Intensities Section 4.3
Midbrain	71.1 (69.5 - 72.7)	63.8 (62.1 - 65.5)	64.6 (62.8 - 66.3)	71.5 (69.8 - 73.1)	51.5 (49.6 - 53.3)	<b>75.4 (73.8 - 77.0)</b>	69.6 (67.9 - 71.4)
Hippocampi	<b>69.8 (68.2 - 71.3)</b>	46.8 (45.2 - 48.4)	64.6 (62.9 - 66.2)	66.1 (64.5 - 67.6)	36.1 (34.5 - 37.6)	63.8 (62.2 - 65.5)	4.2 (3.6 - 4.8)
Basal Ganglia	56.8 (55.0 - 58.5)	51.9 (50.1 - 53.6)	53.3 (51.6 - 55.0)	58.9 (57.2 - 60.6)	56.8 (55.1 - 58.5)	<b>60.3 (58.6 - 62.0)</b>	50.1 (48.3 - 52.0)
Centrum Semiovale	50.6 (49.3 - 52.0)	42.0 (40.7 - 43.4)	48.8 (47.5 - 50.2)	<b>53.0 (51.6 - 54.3)</b>	35.0 (33.9 - 36.1)	49.0 (47.7 - 50.3)	5.7 (5.2 - 6.3)
Average	62.1 +/- 8.7	51.1 +/- 8.1	57.8 +/- 6.9	62.4 +/- 7.0	44.8 +/- 9.5	62.1 +/- 9.4	32.4 +/- 28.3

**Table 3**  
Average number of false positives per scan in the brain datasets. Best performances are indicated in bold.

	GP-Unet (this paper)	GP-Unet no residual Dubost et al. (2017)	Gated Attention Schlemper et al. (2018)	Grad-CAM Selvaraju et al. (2017)	Grad Simonyan et al. (2014)	Guided-backprop Springenberg et al. (2015)	Intensities <a href="#">Section 4.3</a>
Midbrain	<b>1.03 (0.99 - 1.07)</b>	1.19 (1.15 - 1.24)	<b>1.04 (0.99 - 1.09)</b>	1.10 (1.05 - 1.15)	1.40 (1.34 - 1.45)	<b>0.99 (0.94 - 1.03)</b>	1.11 (1.06 - 1.15)
Hippocampi	<b>1.12 (1.06 - 1.17)</b>	1.96 (1.88 - 2.03)	<b>1.13 (1.06 - 1.19)</b>	<b>1.16 (1.10 - 1.22)</b>	2.16 (2.06 - 2.25)	1.23 (1.16 - 1.29)	3.34 (3.22 - 3.45)
Basal Ganglia	<b>1.95 (1.88 - 2.01)</b>	2.33 (2.27 - 2.39)	2.16 (2.10 - 2.23)	2.02 (1.95 - 2.09)	2.06 (1.98 - 2.13)	<b>1.98 (1.91 - 2.04)</b>	2.28 (2.21 - 2.35)
Centrum Semiovale	<b>5.24 (5.04 - 5.43)</b>	6.66 (6.46 - 6.86)	6.23 (6.02 - 6.44)	5.63 (5.44 - 5.82)	7.30 (7.03 - 7.57)	5.92 (5.71 - 6.12)	9.91 (9.62 - 10.21)
Average	2.33 +/- 1.71	3.04 +/- 2.13	2.64 +/- 2.12	2.48 +/- 1.86	3.23 +/- 2.37	2.53 +/- 1.99	4.16 +/- 3.41

**Table 4**  
Average number of false negatives per scan in the brain datasets. Best performances are indicated in bold.

	GP-Unet (this paper)	GP-Unet no residual Dubost et al. (2017)	Gated Attention Schlemper et al. (2018)	Grad-CAM Selvaraju et al. (2017)	Grad Simonyan et al. (2014)	Guided-backprop Springenberg et al. (2015)	Intensities <a href="#">Section 4.3</a>
Midbrain	0.77 (0.71 - 0.83)	0.98 (0.91 - 1.05)	0.94 (0.87 - 1.00)	0.77 (0.71 - 0.82)	1.06 (1.00 - 1.12)	<b>0.65 (0.60 - 0.71)</b>	0.77 (0.72 - 0.83)
Hippocampi	<b>1.14 (1.07 - 1.22)</b>	2.12 (2.01 - 2.23)	1.33 (1.25 - 1.41)	1.32 (1.24 - 1.41)	2.32 (2.21 - 2.43)	1.39 (1.31 - 1.47)	3.50 (3.36 - 3.64)
Basal Ganglia	2.00 (1.85 - 2.14)	2.11 (1.97 - 2.25)	2.08 (1.94 - 2.21)	<b>1.92 (1.78 - 2.06)</b>	1.96 (1.82 - 2.09)	<b>1.88 (1.74 - 2.01)</b>	2.18 (2.03 - 2.33)
Centrum Semiovale	5.83 (5.50 - 6.17)	6.67 (6.30 - 7.03)	5.98 (5.64 - 6.32)	<b>5.63 (5.30 - 5.96)</b>	7.30 (6.92 - 7.68)	5.92 (5.58 - 6.26)	9.91 (9.44 - 10.38)
Average	2.44 +/- 2.01	2.97 +/- 2.18	2.58 +/- 2.00	2.41 +/- 1.90	3.16 +/- 2.43	2.46 +/- 2.04	4.09 +/- 3.50

sensus for the evaluation, which may reduce the risk of mislabeling. We preferred to evaluate the detection using more scans to better encompass the anatomical variability, and we quantified the performance of the single rater by computing her intra-rater agreement on a smaller set.

In our preliminary work on PVS detection in the basal ganglia using GP-Unet no residual (Dubost et al., 2017) we obtained slightly different results than what is presented in the current work. This reflects differences in the test data set, the annotations, method and postprocessing. Our previous annotations (Dubost et al., 2017) were done directly on the segmented and cropped basal ganglia, while the annotations of the current work were done on the full scan. The rater sometimes annotated lesions at the borders of the basal ganglia which are barely visible after preprocessing. In addition, the current work also includes scans without annotations (because the rater found no lesion), where there could have been errors in finding the slice evaluated by the rater. In the current work, Grad reaches better results than in Dubost et al. (2017), because it benefits from the more sophisticated postprocessing: the non-maximum suppression clears the noise in the attention maps.

Next to the methods presented in this paper, we experimented with the perturbation method with masks proposed by Petsiuk et al. (2018). For this method, masks are first sampled in a low dimensional space and resized to the size of the input image. It appeared that the size of this lower dimensional needs to be adapted to the size of the target object in the image. If the target objects are small, one may need to sample relatively large masks. In our experiments, we experimented with a range of values for the size of this low dimensional space, and did not manage to compute discriminative attention maps for PVS, that are small objects relatively to the image resolution.

The work presented in this article implies that pixel-level annotations may not be needed to train accurate models for detection problems. This is especially relevant in medical imaging, where annotation requires expert knowledge and high quality annotations are therefore difficult to obtain. Weakly supervised methods enable learning from large databases, such as UK biobank (Sudlow et al., 2015) or Framingham study (Maillard et al., 2016), with less annotation effort, and could also help to reduce the dependence on annotator biases. The global label may even be more reliable, because for some abnormalities raters can agree well on the presence or global burden of the abnormalities but poorly on their boundaries or spatial distribution.

The variety of challenges present in the brain datasets are well suited to the evaluation of weakly-supervised detection methods. Observations and results might generalize to the detection of other types of small objects, such as microinfarcts, microbleeds, or small white matter hyperintensities.

## 6. Conclusion

We proposed a new weakly supervised detection method, GP-Unet, that uses an encoder-decoder architecture optimized only with global labels such as the count of lesions in a brain region. The decoder part upsamples feature maps and enables the

computation of attention maps at the resolution of the input image, which thus helps the detection of small objects. We also showed the advantage of using regression objectives over classification objectives for the optimization of weakly supervised detection methods, when the target object appears multiple times in the image. We compared the proposed method to four state-of-the-art methods on the detection of digits in MNIST-based datasets, and on the detection of enlarged perivascular spaces – a type of brain lesion – from 3D brain MRI. The best weakly supervised detection methods were Guided-backpropagation (Springenberg et al., 2015), and the proposed method GP-Unet. We noticed that methods based on the gradient of the output of the network, such as Guided-backpropagation, worked best in datasets where the target objects are also the most salient objects. In other datasets, methods using class activation maps, such as GP-Unet, worked best. The performance of the detection enlarged perivascular spaces using the weakly supervised methods was close to the intrarater agreement of an expert rater. The proposed method could consequently facilitate studies of enlarged perivascular and help advance research in their etiology and relationship with cerebrovascular diseases.

## Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## CRedit authorship contribution statement

**Florian Dubost:** Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Validation, Visualization, Software, Writing - original draft, Writing - review & editing. **Hieab Adams:** Conceptualization, Data curation, Funding acquisition, Resources, Writing - review & editing. **Pinar Yilmaz:** Data curation, Investigation, Visualization, Writing - review & editing. **Gerda Bortsova:** Investigation, Methodology, Writing - review & editing. **Gijs van Tulder:** Investigation, Methodology, Writing - review & editing. **M. Arfan Ikram:** Funding acquisition, Resources, Project administration, Writing - review & editing. **Wiro Niessen:** Funding acquisition, Methodology, Project administration, Supervision, Writing - review & editing. **Meike W. Vernooij:** Funding acquisition, Resources, Project administration, Supervision, Writing - review & editing. **Marleen de Bruijne:** Funding acquisition, Methodology, Project administration, Supervision, Writing - review & editing.

## Acknowledgements

This research was funded by The [Netherlands Organisation for Health Research and Development](#) (ZonMw) Project 104003005, with additional support of Netherlands Organisation for Scientific Research (NWO), project NWO-EW VIDI 639.022.010 and project NWO-TTW Perspectief Programme P15-26. This work was partly carried out on the Dutch national e-infrastructure with the support of SURF Cooperative.



Appendix A. Results MNIST – Regression Objectives

Table A.5

**FAUCs MNIST regression.** Each row corresponds to the detection of a different digit. 95 percent confidence interval is indicated in brackets. The average and standard deviation of the performance of each method across all digits is given in the last row. Best performance are indicated in bold.

	GP-Unet (this paper)	GP-Unet no residual Dubost et al. (2017)	Gated Attention Schlemper et al. (2018)	Grad-CAM Selvaraju et al. (2017)	Grad Simonyan et al. (2014)	Guided-backprop Springenberg et al. (2015)
0	97.4 (96.4 - 98.4)	<b>99.7 (99.5 - 99.8)</b>	90.5 (88.9 - 92.0)	89.2 (87.4 - 90.9)	76.1 (73.8 - 78.3)	85.8 (83.8 - 87.5)
1	94.4 (93.5 - 95.1)	<b>99.5 (99.3 - 99.7)</b>	69.7 (67.4 - 72.0)	89.1 (87.8 - 90.4)	66.3 (64.0 - 68.5)	88.3 (86.9 - 89.5)
2	91.7 (90.5 - 92.8)	<b>99.3 (99.0 - 99.5)</b>	87.5 (85.8 - 89.3)	87.8 (85.8 - 89.6)	61.0 (58.5 - 63.6)	81.3 (79.0 - 83.3)
3	97.3 (96.6 - 97.9)	<b>99.6 (99.2 - 99.9)</b>	73.4 (70.9 - 75.6)	91.2 (89.7 - 92.6)	69.8 (67.4 - 72.0)	82.1 (80.0 - 84.1)
4	97.8 (97.2 - 98.3)	<b>99.0 (98.5 - 99.5)</b>	83.9 (81.8 - 86.0)	95.0 (93.8 - 96.0)	86.4 (84.8 - 87.9)	93.5 (92.1 - 94.6)
5	97.1 (96.3 - 97.8)	<b>98.9 (98.4 - 99.4)</b>	79.6 (77.4 - 81.8)	93.6 (92.4 - 94.7)	81.5 (79.7 - 83.3)	88.9 (87.2 - 90.5)
6	98.6 (98.2 - 99.0)	<b>99.9 (99.8 - 99.9)</b>	86.5 (84.4 - 88.5)	98.4 (97.8 - 98.8)	89.3 (87.9 - 90.6)	97.3 (96.6 - 97.9)
7	89.3 (87.6 - 91.0)	<b>99.1 (98.5 - 99.6)</b>	43.9 (41.2 - 46.5)	88.5 (86.9 - 90.0)	69.6 (67.3 - 71.9)	82.4 (80.7 - 84.0)
8	<b>98.8 (98.3 - 99.2)</b>	<b>99.5 (99.1 - 99.8)</b>	73.6 (70.8 - 76.2)	98.2 (97.7 - 98.7)	81.3 (79.1 - 83.3)	95.2 (94.5 - 96.0)
9	<b>97.6 (96.8 - 98.3)</b>	<b>98.6 (98.1 - 99.1)</b>	24.8 (22.6 - 27.1)	94.3 (93.0 - 95.5)	77.9 (75.8 - 80.0)	91.8 (90.4 - 93.2)
Average	96.0 +/- 3.0	99.3 +/- 0.4	71.3 +/- 20.0	92.5 +/- 3.7	75.9 +/- 8.6	88.7 +/- 5.4

Table A.6

**Sensitivity MNIST regression.** Each row corresponds to the detection of a different digit. 95 percent confidence interval is indicated in brackets. The average and standard deviation of the performance of each method across all digits is given in the last row. Best performance are indicated in bold.

	GP-Unet (this paper)	GP-Unet no residual Dubost et al. (2017)	Gated Attention Schlemper et al. (2018)	Grad-CAM Selvaraju et al. (2017)	Grad Simonyan et al. (2014)	Guided-backprop Springenberg et al. (2015)
0	92.7 (91.4 - 93.9)	<b>98.4 (97.8 - 99.1)</b>	75.1 (73.1 - 77.0)	69.2 (67.2 - 71.3)	57.3 (55.3 - 59.2)	63.8 (61.7 - 65.8)
1	78.9 (77.3 - 80.4)	<b>98.3 (97.8 - 98.8)</b>	51.4 (49.4 - 53.4)	70.3 (68.6 - 72.1)	49.5 (47.7 - 51.2)	64.2 (62.5 - 65.8)
2	80.0 (78.2 - 81.8)	<b>96.7 (95.9 - 97.5)</b>	73.5 (71.5 - 75.5)	75.3 (73.4 - 77.2)	44.0 (41.9 - 46.1)	62.8 (60.9 - 64.7)
3	90.1 (88.8 - 91.5)	<b>97.9 (97.4 - 98.4)</b>	55.6 (53.5 - 57.7)	75.0 (73.1 - 76.9)	50.9 (48.9 - 52.9)	57.7 (55.5 - 59.9)
4	90.7 (89.3 - 92.1)	<b>97.0 (96.3 - 97.8)</b>	73.9 (71.8 - 76.0)	84.6 (83.0 - 86.3)	67.9 (66.0 - 69.7)	83.1 (81.4 - 84.8)
5	88.7 (87.2 - 90.2)	<b>96.1 (95.2 - 97.0)</b>	60.0 (57.8 - 62.2)	76.7 (74.6 - 78.7)	58.4 (56.2 - 60.5)	67.3 (65.2 - 69.4)
6	92.2 (91.0 - 93.5)	<b>98.3 (97.7 - 98.9)</b>	73.6 (71.5 - 75.6)	92.4 (91.1 - 93.7)	68.9 (67.1 - 70.7)	84.6 (82.9 - 86.2)
7	76.3 (74.6 - 78.1)	<b>95.9 (94.9 - 97.0)</b>	32.7 (30.7 - 34.6)	73.2 (71.4 - 75.0)	52.7 (50.7 - 54.6)	63.0 (61.2 - 64.8)
8	95.8 (95.0 - 96.5)	<b>98.5 (98.0 - 98.9)</b>	57.7 (55.5 - 59.9)	93.5 (92.5 - 94.4)	60.5 (58.7 - 62.4)	75.0 (73.3 - 76.7)
9	92.3 (91.1 - 93.5)	<b>95.8 (95.0 - 96.6)</b>	17.8 (16.2 - 19.3)	87.9 (86.5 - 89.3)	60.6 (58.7 - 62.5)	78.3 (76.6 - 80.0)
Average	87.8 +/- 6.4	97.3 +/- 1.0	57.1 +/- 18.2	79.8 +/- 8.6	57.1 +/- 7.5	70.0 +/- 9.0




Fig. A.12. FROC MNIST regression. Each subplot corresponds to the detection of a different digit.

**Table A.7**

**FPavg MNIST regression.** Each row corresponds to the detection of a different digit. 95 percent confidence interval is indicated in brackets. The average and standard deviation of the performance of each method across all digits is given in the last row. Best performance are indicated in bold.

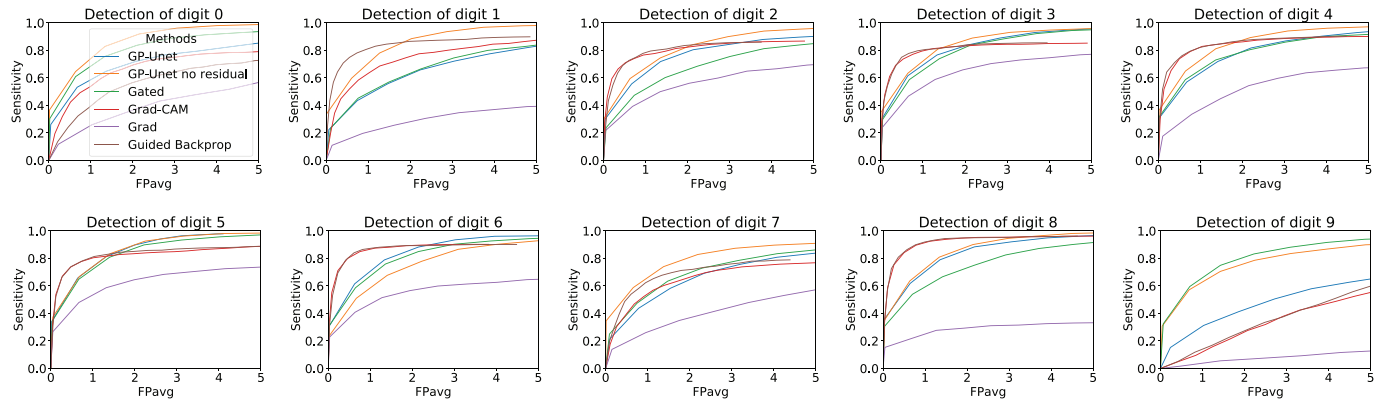
	GP-Unet (this paper)	GP-Unet no residual Dubost et al. (2017)	Gated Attention Schlemper et al. (2018)	Grad-CAM Selvaraju et al. (2017)	Grad Simonyan et al. (2014)	Guided-backprop Springenberg et al. (2015)
0	0.07 (0.05 - 0.08)	<b>0.02 (0.01 - 0.03)</b>	0.14 (0.11 - 0.17)	0.16 (0.13 - 0.18)	0.21 (0.18 - 0.24)	0.27 (0.23 - 0.31)
1	0.14 (0.11 - 0.16)	<b>0.02 (0.01 - 0.03)</b>	0.42 (0.36 - 0.47)	0.34 (0.30 - 0.38)	0.48 (0.42 - 0.53)	0.36 (0.31 - 0.40)
2	0.29 (0.25 - 0.32)	<b>0.03 (0.01 - 0.04)</b>	0.32 (0.28 - 0.36)	0.37 (0.33 - 0.41)	0.79 (0.73 - 0.86)	0.42 (0.38 - 0.47)
3	0.11 (0.09 - 0.14)	<b>0.00 (0.00 - 0.00)</b>	0.47 (0.42 - 0.52)	0.26 (0.22 - 0.29)	0.52 (0.46 - 0.58)	0.43 (0.38 - 0.48)
4	0.13 (0.10 - 0.15)	<b>0.03 (0.02 - 0.04)</b>	0.39 (0.34 - 0.43)	0.18 (0.15 - 0.21)	0.19 (0.16 - 0.22)	0.19 (0.16 - 0.22)
5	0.13 (0.11 - 0.16)	<b>0.03 (0.02 - 0.05)</b>	0.36 (0.31 - 0.40)	0.18 (0.15 - 0.21)	0.30 (0.25 - 0.34)	0.23 (0.20 - 0.26)
6	0.04 (0.03 - 0.06)	<b>0.01 (0.00 - 0.01)</b>	0.40 (0.35 - 0.44)	0.10 (0.08 - 0.12)	0.14 (0.11 - 0.17)	0.06 (0.05 - 0.08)
7	0.26 (0.23 - 0.29)	<b>0.02 (0.01 - 0.03)</b>	1.02 (0.93 - 1.11)	0.36 (0.32 - 0.41)	0.57 (0.51 - 0.62)	0.41 (0.36 - 0.45)
8	0.08 (0.06 - 0.10)	<b>0.02 (0.01 - 0.03)</b>	0.44 (0.38 - 0.49)	0.11 (0.09 - 0.14)	0.21 (0.17 - 0.25)	0.07 (0.05 - 0.09)
9	0.16 (0.13 - 0.18)	<b>0.04 (0.03 - 0.06)</b>	1.19 (1.07 - 1.30)	0.25 (0.22 - 0.28)	0.33 (0.29 - 0.38)	0.19 (0.16 - 0.23)
Average	0.14 +/- 0.08	0.02 +/- 0.01	0.51 +/- 0.31	0.23 +/- 0.10	0.37 +/- 0.20	0.26 +/- 0.13

**Table A.8**

**FNavg MNIST regression.** Each row corresponds to the detection of a different digit. 95 percent confidence interval is indicated in brackets. The average and standard deviation of the performance of each method across all digits is given in the last row. Best performance are indicated in bold.

	GP-Unet (this paper)	GP-Unet no residual Dubost et al. (2017)	Gated Attention Schlemper et al. (2018)	Grad-CAM Selvaraju et al. (2017)	Grad Simonyan et al. (2014)	Guided-backprop Springenberg et al. (2015)
0	0.11 (0.08 - 0.14)	<b>0.02 (0.01 - 0.03)</b>	0.37 (0.32 - 0.43)	0.47 (0.41 - 0.53)	0.62 (0.55 - 0.69)	0.54 (0.47 - 0.61)
1	0.35 (0.30 - 0.40)	<b>0.03 (0.02 - 0.05)</b>	0.75 (0.67 - 0.82)	0.49 (0.43 - 0.56)	0.82 (0.73 - 0.91)	0.58 (0.52 - 0.64)
2	0.30 (0.25 - 0.34)	<b>0.05 (0.03 - 0.07)</b>	0.40 (0.34 - 0.45)	0.37 (0.32 - 0.43)	0.82 (0.73 - 0.90)	0.56 (0.50 - 0.63)
3	0.15 (0.12 - 0.18)	<b>0.04 (0.02 - 0.05)</b>	0.58 (0.52 - 0.65)	0.34 (0.30 - 0.39)	0.68 (0.61 - 0.76)	0.56 (0.49 - 0.62)
4	0.12 (0.10 - 0.15)	<b>0.05 (0.03 - 0.06)</b>	0.36 (0.31 - 0.41)	0.22 (0.18 - 0.25)	0.49 (0.43 - 0.56)	0.25 (0.21 - 0.29)
5	0.14 (0.11 - 0.16)	<b>0.05 (0.04 - 0.07)</b>	0.50 (0.44 - 0.56)	0.29 (0.25 - 0.34)	0.54 (0.48 - 0.60)	0.42 (0.37 - 0.47)
6	0.11 (0.08 - 0.13)	<b>0.02 (0.01 - 0.03)</b>	0.35 (0.31 - 0.40)	0.09 (0.07 - 0.12)	0.47 (0.41 - 0.52)	0.22 (0.18 - 0.26)
7	0.36 (0.31 - 0.41)	<b>0.04 (0.03 - 0.06)</b>	1.00 (0.90 - 1.10)	0.43 (0.37 - 0.49)	0.73 (0.65 - 0.81)	0.58 (0.51 - 0.65)
8	0.07 (0.05 - 0.09)	<b>0.02 (0.01 - 0.04)</b>	0.58 (0.50 - 0.65)	0.11 (0.08 - 0.13)	0.59 (0.52 - 0.66)	0.37 (0.32 - 0.43)
9	0.11 (0.09 - 0.14)	<b>0.07 (0.05 - 0.09)</b>	1.22 (1.09 - 1.34)	0.20 (0.16 - 0.24)	0.63 (0.55 - 0.70)	0.36 (0.31 - 0.41)
Average	0.18 +/- 0.10	0.04 +/- 0.02	0.61 +/- 0.28	0.30 +/- 0.14	0.64 +/- 0.12	0.44 +/- 0.13

**Appendix B. Results MNIST – Classification Objectives**



**Fig. B.13. FROC MNIST classification.** Each subplot corresponds to the detection of a different digit.

**Table B.9**

**FAUCs MNIST classification.** Each row corresponds to the detection of a different digit. 95 percent confidence interval is indicated in brackets. The average and standard deviation of the performance of each method across all digits is given in the last row. Best performance are indicated in bold.

	GP-Unet (this paper)	GP-Unet no residual Dubost et al. (2017)	Gated Attention Schlemper et al. (2018)	Grad-CAM Selvaraju et al. (2017)	Grad Simonyan et al. (2014)	Guided-backprop Springenberg et al. (2015)
0	69.3 (66.7 - 71.7)	<b>85.8 (84.6 - 87.1)</b>	79.5 (77.4 - 81.6)	65.2 (61.8 - 68.6)	37.5 (34.4 - 40.6)	54.1 (50.7 - 57.7)
1	62.6 (60.3 - 65.0)	<b>83.9 (82.5 - 85.3)</b>	63.8 (61.3 - 66.3)	72.3 (69.5 - 75.3)	28.6 (25.6 - 31.7)	<b>83.2 (80.3 - 86.1)</b>
2	75.5 (73.4 - 77.5)	<b>80.3 (78.7 - 81.9)</b>	65.9 (63.7 - 68.3)	<b>80.3 (76.9 - 83.6)</b>	54.9 (50.7 - 58.6)	<b>80.1 (76.5 - 83.5)</b>
3	80.6 (79.0 - 82.2)	<b>83.4 (81.7 - 84.9)</b>	78.9 (77.0 - 80.7)	<b>80.7 (77.0 - 84.3)</b>	63.7 (59.9 - 67.4)	<b>80.9 (77.2 - 84.5)</b>
4	77.1 (74.9 - 79.1)	<b>84.1 (82.6 - 85.7)</b>	77.1 (75.2 - 79.0)	<b>83.7 (80.7 - 86.5)</b>	51.3 (47.4 - 55.0)	<b>84.3 (81.1 - 87.3)</b>
5	<b>85.5 (83.9 - 86.8)</b>	<b>85.5 (84.3 - 86.7)</b>	<b>83.4 (81.9 - 84.9)</b>	<b>81.7 (78.5 - 84.7)</b>	62.0 (58.3 - 65.8)	<b>82.5 (79.0 - 85.8)</b>
6	82.8 (81.1 - 84.3)	74.1 (71.9 - 76.2)	80.0 (77.8 - 81.9)	<b>86.6 (83.7 - 89.3)</b>	54.6 (50.4 - 58.9)	<b>86.4 (83.6 - 89.2)</b>
7	63.7 (61.1 - 66.2)	<b>77.9 (75.9 - 79.9)</b>	67.7 (65.4 - 70.0)	63.5 (60.1 - 67.0)	38.4 (35.8 - 41.2)	68.4 (64.7 - 72.1)
8	82.7 (81.0 - 84.2)	84.7 (83.2 - 86.0)	73.2 (71.1 - 75.3)	<b>91.6 (89.4 - 93.6)</b>	28.6 (24.1 - 33.1)	<b>91.9 (89.9 - 93.7)</b>
9	44.3 (41.5 - 47.3)	74.3 (72.2 - 76.5)	<b>78.1 (76.0 - 80.0)</b>	30.4 (27.6 - 33.3)	7.2 (5.7 - 8.8)	31.7 (28.8 - 34.7)
Average	72.4 +/- 12.0	81.4 +/- 4.3	74.8 +/- 6.4	73.6 +/- 16.7	42.7 +/- 17.0	74.4 +/- 17.4

**Table B.10**

**Sensitivity MNIST classification.** Each row corresponds to the detection of a different digit. 95 percent confidence interval is indicated in brackets. The average and standard deviation of the performance of each method across all digits is given in the last row. Best performance are indicated in bold.

	GP-Unet (this paper)	GP-Unet no residual Dubost et al. (2017)	Gated Attention Schlemper et al. (2018)	Grad-CAM Selvaraju et al. (2017)	Grad Simonyan et al. (2014)	Guided-backprop Springenberg et al. (2015)
1	26.0 (24.7 - 27.4)	<b>36.7 (34.9 - 38.5)</b>	30.3 (28.7 - 31.9)	19.4 (18.0 - 20.8)	11.8 (10.6 - 12.9)	13.6 (12.3 - 14.9)
2	22.3 (21.0 - 23.6)	<b>33.5 (31.9 - 35.1)</b>	23.4 (22.1 - 24.7)	19.4 (18.2 - 20.6)	11.0 (9.8 - 12.2)	25.9 (24.6 - 27.3)
3	<b>30.1 (28.6 - 31.7)</b>	<b>30.4 (28.8 - 31.9)</b>	23.7 (22.3 - 25.1)	27.3 (25.8 - 28.8)	21.6 (20.0 - 23.3)	25.0 (23.5 - 26.5)
4	30.3 (28.8 - 31.7)	<b>35.1 (33.5 - 36.8)</b>	29.8 (28.4 - 31.2)	28.2 (26.7 - 29.6)	23.9 (22.3 - 25.6)	26.9 (25.5 - 28.3)
5	30.8 (29.2 - 32.4)	<b>36.5 (34.8 - 38.3)</b>	31.2 (29.6 - 32.8)	29.6 (28.1 - 31.1)	17.6 (15.9 - 19.3)	30.8 (29.2 - 32.3)
6	34.1 (32.5 - 35.8)	<b>37.4 (35.6 - 39.1)</b>	34.4 (32.8 - 36.1)	31.4 (29.8 - 33.0)	26.7 (25.0 - 28.4)	30.1 (28.5 - 31.6)
7	<b>30.9 (29.4 - 32.3)</b>	23.9 (22.7 - 25.0)	<b>31.0 (29.6 - 32.5)</b>	<b>31.3 (29.7 - 32.8)</b>	22.9 (21.1 - 24.7)	<b>30.8 (29.3 - 32.4)</b>
8	21.6 (20.3 - 22.8)	<b>34.6 (32.9 - 36.3)</b>	25.0 (23.5 - 26.5)	17.5 (16.2 - 18.8)	13.9 (12.7 - 15.0)	19.1 (17.9 - 20.3)
9	<b>34.7 (33.1 - 36.4)</b>	<b>34.0 (32.4 - 35.6)</b>	30.7 (29.2 - 32.2)	<b>33.5 (31.9 - 35.1)</b>	15.3 (13.3 - 17.3)	<b>33.7 (32.1 - 35.3)</b>
Average	15.2 (13.9 - 16.6)	<b>30.7 (29.2 - 32.2)</b>	<b>32.4 (30.8 - 34.0)</b>	4.9 (4.1 - 5.7)	1.7 (1.1 - 2.2)	5.1 (4.3 - 6.0)
Average	27.6 +/- 5.9	33.3 +/- 3.9	29.2 +/- 3.6	24.2 +/- 8.4	16.6 +/- 7.1	24.1 +/- 8.5

**Table B.11**

**FPavg MNIST classification.** Each row corresponds to the detection of a different digit. 95 percent confidence interval is indicated in brackets. The average and standard deviation of the performance of each method across all digits is given in the last row. Best performance are indicated in bold.

	GP-Unet (this paper)	GP-Unet no residual Dubost et al. (2017)	Gated Attention Schlemper et al. (2018)	Grad-CAM Selvaraju et al. (2017)	Grad Simonyan et al. (2014)	Guided-backprop Springenberg et al. (2015)
0	0.05 (0.03 - 0.06)	<b>0.02 (0.01 - 0.03)</b>	<b>0.02 (0.01 - 0.04)</b>	0.15 (0.12 - 0.17)	0.23 (0.20 - 0.26)	0.21 (0.18 - 0.24)
1	0.06 (0.05 - 0.08)	<b>0.00 (0.00 - 0.01)</b>	0.11 (0.08 - 0.13)	0.10 (0.08 - 0.13)	0.14 (0.12 - 0.17)	0.04 (0.02 - 0.05)
2	0.04 (0.03 - 0.05)	<b>0.01 (0.00 - 0.02)</b>	0.06 (0.04 - 0.08)	<b>0.03 (0.01 - 0.04)</b>	0.05 (0.03 - 0.06)	0.05 (0.03 - 0.07)
3	0.02 (0.01 - 0.04)	<b>0.00 (0.00 - 0.01)</b>	0.04 (0.03 - 0.05)	0.02 (0.01 - 0.03)	0.03 (0.02 - 0.04)	0.03 (0.02 - 0.04)
4	<b>0.02 (0.01 - 0.03)</b>	<b>0.02 (0.01 - 0.03)</b>	<b>0.01 (0.00 - 0.02)</b>	0.04 (0.03 - 0.06)	0.10 (0.08 - 0.13)	0.03 (0.02 - 0.05)
5	<b>0.02 (0.01 - 0.03)</b>	0.05 (0.03 - 0.06)	<b>0.03 (0.02 - 0.04)</b>	0.04 (0.03 - 0.05)	0.05 (0.03 - 0.06)	0.05 (0.03 - 0.06)
6	<b>0.03 (0.01 - 0.04)</b>	<b>0.03 (0.02 - 0.04)</b>	<b>0.02 (0.01 - 0.03)</b>	<b>0.02 (0.01 - 0.03)</b>	0.03 (0.02 - 0.05)	0.03 (0.02 - 0.04)
7	0.10 (0.08 - 0.12)	<b>0.01 (0.00 - 0.02)</b>	0.09 (0.07 - 0.11)	0.11 (0.08 - 0.13)	0.15 (0.12 - 0.18)	0.08 (0.06 - 0.10)
8	<b>0.01 (0.00 - 0.01)</b>	<b>0.01 (0.00 - 0.02)</b>	0.04 (0.02 - 0.05)	0.04 (0.02 - 0.05)	0.04 (0.03 - 0.06)	0.03 (0.02 - 0.04)
9	0.24 (0.20 - 0.27)	<b>0.04 (0.03 - 0.05)</b>	0.07 (0.05 - 0.08)	0.42 (0.38 - 0.45)	0.48 (0.44 - 0.51)	0.42 (0.38 - 0.45)
Average	0.06 +/- 0.06	0.02 +/- 0.01	0.05 +/- 0.03	0.10 +/- 0.11	0.13 +/- 0.13	0.10 +/- 0.12

**Table B.12**

**FNavg MNIST classification.** Each row corresponds to the detection of a different digit. 95 percent confidence interval is indicated in brackets. The average and standard deviation of the performance of each method across all digits is given in the last row. Best performance are indicated in bold.

	GP-Unet (this paper)	GP-Unet no residual Dubost et al. (2017)	Gated Attention Schlemper et al. (2018)	Grad-CAM Selvaraju et al. (2017)	Grad Simonyan et al. (2014)	Guided-backprop Springenberg et al. (2015)
0	1.13 (1.01 - 1.25)	<b>1.07 (0.95 - 1.19)</b>	1.11 (0.99 - 1.23)	1.22 (1.09 - 1.35)	1.31 (1.17 - 1.44)	1.28 (1.15 - 1.41)
1	1.34 (1.20 - 1.48)	<b>1.24 (1.11 - 1.38)</b>	1.32 (1.18 - 1.46)	1.36 (1.22 - 1.50)	1.52 (1.36 - 1.68)	1.29 (1.15 - 1.43)
2	<b>1.15 (1.02 - 1.28)</b>	<b>1.14 (1.02 - 1.27)</b>	1.20 (1.07 - 1.32)	1.17 (1.04 - 1.30)	1.27 (1.13 - 1.42)	1.19 (1.06 - 1.32)
3	1.05 (0.93 - 1.16)	<b>1.02 (0.90 - 1.14)</b>	1.05 (0.93 - 1.16)	1.06 (0.95 - 1.18)	1.14 (1.02 - 1.27)	1.08 (0.96 - 1.19)
4	1.09 (0.97 - 1.21)	<b>1.05 (0.93 - 1.17)</b>	1.08 (0.96 - 1.21)	1.09 (0.97 - 1.22)	1.27 (1.13 - 1.42)	1.09 (0.97 - 1.21)
5	0.99 (0.88 - 1.10)	<b>0.97 (0.86 - 1.08)</b>	0.98 (0.87 - 1.10)	1.00 (0.89 - 1.12)	1.09 (0.96 - 1.21)	1.01 (0.90 - 1.12)
6	<b>1.11 (0.99 - 1.23)</b>	1.16 (1.04 - 1.28)	<b>1.11 (0.99 - 1.23)</b>	<b>1.11 (0.99 - 1.23)</b>	1.26 (1.12 - 1.40)	<b>1.12 (1.00 - 1.24)</b>
7	1.26 (1.13 - 1.40)	<b>1.17 (1.04 - 1.31)</b>	1.24 (1.11 - 1.38)	1.32 (1.18 - 1.45)	1.37 (1.23 - 1.51)	1.29 (1.16 - 1.43)
8	<b>1.06 (0.94 - 1.18)</b>	<b>1.06 (0.94 - 1.19)</b>	1.09 (0.96 - 1.21)	<b>1.07 (0.94 - 1.19)</b>	1.34 (1.18 - 1.49)	<b>1.06 (0.94 - 1.19)</b>
9	1.32 (1.18 - 1.46)	<b>1.15 (1.02 - 1.29)</b>	<b>1.15 (1.01 - 1.28)</b>	1.46 (1.31 - 1.61)	1.52 (1.36 - 1.68)	1.46 (1.31 - 1.61)
Average	1.15 +/- 0.11	1.10 +/- 0.08	1.13 +/- 0.09	1.19 +/- 0.14	1.31 +/- 0.13	1.19 +/- 0.13

Appendix C. Results MNIST – GP-Unet architectures




Fig. C.14. FROC MNIST architecture. Each subplot corresponds to the detection of a different digit. GP-Unet is the standard GP-Unet architecture. In GP-Unet No Skip, blockwise skip connections are removed. In GP-Unet Max Pool, the global average pooling is replaced by global max pooling.

Table C.13

FAUCs MNIST architecture. Each row corresponds to the detection of a different digit. 95 percent confidence interval is indicated in brackets. The average and standard deviation of the performance of each method across all digits is given in the last row. Best performance are indicated in bold.

	GP-Unet	GP-Unet no skip	GP-Unet pax pooling
0	<b>97.4 (96.3 - 98.3)</b>	<b>98.0 (97.5 - 98.5)</b>	89.2 (87.2 - 91.3)
1	<b>94.4 (93.6 - 95.2)</b>	86.9 (85.6 - 88.2)	80.0 (77.2 - 82.6)
2	<b>91.7 (90.4 - 92.9)</b>	<b>92.9 (91.4 - 94.3)</b>	80.3 (77.5 - 83.0)
3	<b>97.3 (96.6 - 97.9)</b>	<b>97.9 (97.4 - 98.5)</b>	88.6 (86.4 - 90.8)
4	<b>97.8 (97.2 - 98.3)</b>	<b>96.2 (95.2 - 97.0)</b>	82.9 (80.2 - 85.5)
5	<b>97.1 (96.4 - 97.8)</b>	<b>97.6 (97.1 - 98.1)</b>	86.1 (83.6 - 88.6)
6	<b>98.6 (98.1 - 99.1)</b>	<b>98.1 (97.5 - 98.7)</b>	87.6 (85.2 - 89.9)
7	<b>89.3 (87.6 - 90.9)</b>	0.8 (0.5 - 1.2)	56.3 (52.8 - 59.8)
8	<b>98.8 (98.2 - 99.2)</b>	<b>98.8 (98.4 - 99.1)</b>	87.7 (85.3 - 89.9)
9	<b>97.6 (96.8 - 98.2)</b>	<b>96.9 (96.1 - 97.6)</b>	78.4 (75.5 - 81.1)
Average	96.0 +/- 3.0	86.4 +/- 28.7	81.7 +/- 9.3

Table C.14

Sensitivity MNIST architecture. Each row corresponds to the detection of a different digit. 95 percent confidence interval is indicated in brackets. The average and standard deviation of the performance of each method across all digits is given in the last row. Best performance are indicated in bold.

	GP-Unet	GP-Unet no skip	GP-Unet pax pooling
0	<b>92.7 (91.4 - 93.9)</b>	<b>90.9 (89.6 - 92.3)</b>	81.8 (80.1 - 83.5)
1	<b>78.9 (77.3 - 80.4)</b>	63.3 (61.5 - 65.0)	71.9 (70.1 - 73.8)
2	<b>80.0 (78.2 - 81.8)</b>	<b>83.4 (81.7 - 85.1)</b>	75.8 (74.0 - 77.7)
3	90.1 (88.8 - 91.5)	<b>91.8 (90.7 - 93.0)</b>	82.8 (81.1 - 84.5)
4	<b>90.7 (89.3 - 92.1)</b>	87.8 (86.3 - 89.2)	76.1 (74.1 - 78.1)
5	<b>88.7 (87.2 - 90.2)</b>	<b>89.4 (87.9 - 90.8)</b>	74.7 (72.7 - 76.8)
6	<b>92.2 (91.0 - 93.5)</b>	<b>91.9 (90.5 - 93.3)</b>	81.4 (79.5 - 83.4)
7	<b>76.3 (74.6 - 78.1)</b>	0.6 (0.3 - 0.9)	55.2 (52.9 - 57.5)
8	<b>95.8 (95.0 - 96.5)</b>	<b>94.5 (93.7 - 95.4)</b>	83.7 (82.1 - 85.3)
9	<b>92.3 (91.1 - 93.5)</b>	90.1 (88.8 - 91.5)	70.8 (68.8 - 72.9)
Average	87.8 +/- 6.4	78.4 +/- 27.3	75.4 +/- 8.0

**Table C.15**


**FPavg MNIST architecture.** Each row corresponds to the detection of a different digit. 95 percent confidence interval is indicated in brackets. The average and standard deviation of the performance of each method across all digits is given in the last row. Best performance are indicated in bold.

	GP-Unet	GP-Unet no skip	GP-Unet pax pooling
0	<b>0.07 (0.05 - 0.08)</b>	<b>0.07 (0.05 - 0.09)</b>	<b>0.08 (0.06 - 0.10)</b>
1	<b>0.14 (0.11 - 0.16)</b>	0.30 (0.26 - 0.35)	<b>0.16 (0.13 - 0.19)</b>
2	0.29 (0.25 - 0.32)	0.22 (0.18 - 0.25)	<b>0.15 (0.12 - 0.17)</b>
3	0.11 (0.09 - 0.14)	<b>0.06 (0.05 - 0.08)</b>	<b>0.09 (0.07 - 0.12)</b>
4	0.13 (0.10 - 0.15)	0.10 (0.08 - 0.13)	<b>0.06 (0.04 - 0.07)</b>
5	0.13 (0.11 - 0.16)	<b>0.10 (0.08 - 0.12)</b>	<b>0.09 (0.06 - 0.11)</b>
6	<b>0.04 (0.03 - 0.06)</b>	0.08 (0.06 - 0.10)	0.11 (0.09 - 0.14)
7	<b>0.26 (0.23 - 0.29)</b>	1.99 (1.98 - 2.00)	<b>0.21 (0.17 - 0.24)</b>
8	0.08 (0.06 - 0.10)	0.06 (0.05 - 0.08)	<b>0.03 (0.02 - 0.04)</b>
9	0.16 (0.13 - 0.18)	0.19 (0.16 - 0.22)	<b>0.10 (0.08 - 0.13)</b>
Average	0.14 +/- 0.08	0.32 +/- 0.56	0.11 +/- 0.05


**Table C.16**

**FNavg MNIST architecture.** Each row corresponds to the detection of a different digit. 95 percent confidence interval is indicated in brackets. The average and standard deviation of the performance of each method across all digits is given in the last row. Best performance are indicated in bold.

	GP-Unet	GP-Unet no skip	GP-Unet pax pooling
0	<b>0.11 (0.08 - 0.14)</b>	<b>0.13 (0.10 - 0.16)</b>	0.30 (0.25 - 0.35)
1	<b>0.35 (0.30 - 0.40)</b>	0.60 (0.53 - 0.67)	0.49 (0.42 - 0.56)
2	0.30 (0.25 - 0.34)	<b>0.25 (0.21 - 0.29)</b>	0.39 (0.34 - 0.45)
3	<b>0.15 (0.12 - 0.18)</b>	<b>0.13 (0.10 - 0.16)</b>	0.26 (0.22 - 0.31)
4	<b>0.12 (0.10 - 0.15)</b>	0.18 (0.14 - 0.21)	0.37 (0.31 - 0.42)
5	<b>0.14 (0.11 - 0.16)</b>	<b>0.13 (0.11 - 0.16)</b>	0.36 (0.30 - 0.41)
6	<b>0.11 (0.08 - 0.13)</b>	<b>0.10 (0.08 - 0.12)</b>	0.26 (0.22 - 0.30)
7	<b>0.36 (0.31 - 0.41)</b>	1.58 (1.42 - 1.74)	0.71 (0.62 - 0.80)
8	<b>0.07 (0.05 - 0.09)</b>	<b>0.09 (0.07 - 0.11)</b>	0.26 (0.21 - 0.31)
9	<b>0.11 (0.09 - 0.14)</b>	0.15 (0.12 - 0.18)	0.47 (0.41 - 0.54)
Average	0.18 +/- 0.10	0.33 +/- 0.44	0.39 +/- 0.13



**Fig. C.15. FAUCs on the MNIST dataset for different architectures of the proposed method: GP-Unet.** Each subplot corresponds to the detection of a different digit. FAUCs are displayed with confidence intervals computed by bootstrapping the test set. A: GP-Unet; B: GP-Unet without blockwise skip connections; C: GP-Unet with Global Max Pooling. For digit 7, GP-Unet without blockwise skip connections convergence to a very high value. We tried repeating the experiments with different random initializations of the weights and let the optimization run longer, but we achieved the same results. This supports the argument that adding blockwise skip connections eases the optimization.



**Fig. C.16. Attention Maps for the different architectures of GP-UNet.** In the first row, the digits 1 had to be detected. In the second row, digits 4 had to be detected. GP-UNet is the standard GP-UNet architecture. In GP-UNet No Skip, blockwise skip connections are removed. In GP-UNet Max Pool, the global average pooling is replaced by global max pooling.

## References

- Adams, H.H., Cavalieri, M., Verhaaren, B.F., Bos, D., van der Lugt, A., Enzinger, C., Vernooij, M.W., Schmidt, R., Ikram, M.A., 2013. Rating method for dilated virchow-robin spaces on magnetic resonance imaging. *Stroke* 44 (6), 1732–1735.
- Adams, H.H., Schwingenschuh, P., van der Lugt, A., Katschnig-Winter, P., Hofman, A., Krestin, G., Vernooij, M., Schmidt, R., Ikram, M.A., 2014. Determinants of enlarged virchow-robin spaces: the universe consortium. *Alzheimer's & Dementia: The Journal of the Alzheimer's Association* 10 (4), P408.
- Adebayo, J., Gilmer, J., Muelly, M., Goodfellow, I., Hardt, M., Kim, B., 2018. Sanity checks for saliency maps. In: *Advances in Neural Information Processing Systems*, pp. 9524–9535.
- Alam, M.M., Islam, M.T., 2019. Machine learning approach of automatic identification and counting of blood cells. *Healthc Technol Lett* 6 (4), 103–108.
- Andrews, S., Tsochantaridis, I., Hofmann, T., 2003. Support vector machines for multiple-instance learning. In: *Advances in neural information processing systems*, pp. 577–584.
- Ballerini, L., Lovreglio, R., Hernández, M.d.C.V., Ramirez, J., MacIntosh, B.J., Black, S.E., Wardlaw, J.M., 2018. Perivascular spaces segmentation in brain mri using optimal 3D filtering. *Sci Rep* 8 (1), 1–11.
- Bandos, A.I., Rockette, H.E., Song, T., Gur, D., 2009. Area under the free-response ROC curve (FROC) and a related summary index. *Biometrics* 65 (1), 247–256.
- Boespflug, E.L., Schwartz, D.L., Lahna, D., Pollock, J., Iliff, J.J., Kaye, J.A., Rooney, W., Silbert, L.C., 2018. Mr imaging-based multimodal autoidentification of perivascular spaces (MMAPS): automated morphologic segmentation of enlarged perivascular spaces at clinical field strength. *Radiology* 286 (2), 632–642.
- Bortsova, G., Dubost, F., Ørting, S., Katramados, I., Hogeweg, L., Thomsen, L., Wille, M., de Bruijne, M., 2018. Deep learning from label proportions for emphysema quantification. In: *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 768–776.
- Brown, R., Benveniste, H., Black, S.E., Charpak, S., Dichgans, M., Joutel, A., Nedergaard, M., Smith, K.J., Zlokovic, B.V., Wardlaw, J.M., 2018. Understanding the role of the perivascular space in cerebral small vessel disease. *Cardiovasc. Res.* 114 (11), 1462–1473.
- Chollet, F., et al., 2015. Keras. <https://github.com/fchollet/keras>.
- Cole, J.H., Poudel, R.P., Tsagkrasoulis, D., Caan, M.W., Steves, C., Spector, T.D., Montana, G., 2017. Predicting brain age with deep learning from raw imaging data results in a reliable and heritable biomarker. *Neuroimage* 163, 115–124.
- Desikan, R.S., Ségonne, 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., Killiany, R.J., 2006. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage* 31 (3), 968–980.
- Dubost, F., Adams, H., Bortsova, G., Ikram, M.A., Niessen, W., Vernooij, M., de Bruijne, M., 2019. 3D regression neural network for the quantification of enlarged perivascular spaces in brain mri. *Med Image Anal* 51, 89–100.
- Dubost, F., Bortsova, G., Adams, H., Ikram, A., Niessen, W.J., Vernooij, M., de Bruijne, M., 2017. GP-UNet: Lesion detection from weak labels with a 3D regression network. In: *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 214–221.
- Dubost, F., Yilmaz, P., Adams, H., Bortsova, G., Ikram, M.A., Niessen, W., Vernooij, M., de Bruijne, M., 2019. Enlarged perivascular spaces in brain MRI: automated quantification in four regions. *Neuroimage* 185, 534–544.
- Duperron, M.-G., Tzourio, C., Schilling, S., Zhu, Y.-C., Soumaré, A., Mazoyer, B., Debette, S., 2019. High dilated perivascular space burden: a new mri marker for risk of intracerebral hemorrhage. *Neurobiol. Aging* 84, 158–165.
- Felzenszwalb, P.F., Girshick, R.B., McAllester, D., Ramanan, D., 2010. Object detection with discriminatively trained part-based models. *IEEE Trans Pattern Anal Mach Intell* 32 (9), 1627–1645.
- Fong, R.C., Vedaldi, A., 2017. Interpretable explanations of black boxes by meaningful perturbation. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Gutierrez, J., DiTullio, M., Cheung, Y.K.K., Alperin, N., Bagci, A., Sacco, R.L., Wright, C.B., Elkind, M.S., Rundek, T., 2019. Brain arterial dilatation modifies the association between extracranial pulsatile hemodynamics and brain perivascular spaces: the northern manhattan study. *Hypertens. Res.* 42 (7), 1019–1028.
- Hwang, S., Kim, H.-E., 2016. Self-transfer learning for weakly supervised lesion localization. In: *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 239–246.
- Ikram, M.A., van der Lugt, A., Niessen, W.J., Koudstaal, P.J., Krestin, G.P., Hofman, A., Bos, D., Vernooij, M.W., 2015. The rotterdam scan study: design update 2016 and main findings. *Eur. J. Epidemiol.* 30 (12), 1299–1315.
- Ilse, M., Tomczak, J.M., Welling, M., 2018. Attention-based deep multiple instance learning. *ICML*.
- Jungreis, C.A., Kanal, E., Hirsch, W., Martinez, A., Moosy, J., 1988. Normal perivascular spaces mimicking lacunar infarction: MR imaging. *Radiology* 169 (1), 101–104.
- Korbar, B., Olofson, A.M., Mirafior, A.P., Nicka, C.M., Suriawinata, M.A., Torresani, L., Suriawinata, A.A., Hassanpour, S., 2017. Looking under the hood: Deep neural network visualization to interpret whole-slide image analysis outcomes for colorectal polyps. In: *Proceedings of the Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*. IEEE, pp. 821–827.
- Kuhn, H.W., 1955. The hungarian method for the assignment problem. *Naval research logistics quarterly* 2 (1–2), 83–97.
- Laina, I., Ruppert, C., Belagiannis, V., Tombari, F., Navab, N., 2016. Deeper depth prediction with fully convolutional residual networks. In: *Proceedings of the 3D Vision (3DV), 2016 Fourth International Conference on*. IEEE, pp. 239–248.
- LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86 (11), 2278–2324.
- Lee, C.-Y., Xie, S., Gallagher, P., Zhang, Z., Tu, Z., 2015. Deeply-supervised nets. In: *Artificial Intelligence and Statistics*, pp. 562–570.
- Lempitsky, V., Zisserman, A., 2010. Learning to count objects in images. In: *Advances in neural information processing systems*, pp. 1324–1332.
- Maillard, P., Mitchell, G.F., Himali, J.J., Beiser, A., Tsao, C.W., Pase, M.P., Satizabal, C.L., Vasani, R.S., Seshadri, S., DeCarli, C., 2016. Effects of arterial stiffness on brain integrity in young adults from the framingham heart study. *Stroke* 47 (4), 1030–1036.
- Mundhenk, T.N., Konjevod, G., Sakla, W.A., Boakye, K., 2016. A large contextual dataset for classification, detection and counting of cars with deep learning. In: *Proceedings of the European Conference on Computer Vision*. Springer, pp. 785–800.
- Oktay, O., Schlemper, J., Folgoc, L.L., Lee, M., Heinrich, M., Misawa, K., Mori, K., McDonagh, S., Hammerla, N.Y., Kainz, B., et al., 2018. Attention U-Net: learning where to look for the pancreas. *Medical Imaging with Deep Learning Conference*.
- Paul Cohen, J., Boucher, G., Glastonbury, C.A., Lo, H.Z., Bengio, Y., 2017. Count-ception: Counting by fully convolutional redundant counting. In: *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pp. 18–26.
- Petsiuk, V., Das, A., Saenko, K., 2018. Rise: Randomized input sampling for explanation of black-box models. In: *Proceedings of the British Machine Vision Conference*.
- Potter, G.M., Doubal, F.N., Jackson, C.A., Chappell, F.M., Sudlow, C.L., Dennis, M.S.,

- Wardlaw, J.M., 2015. Enlarged perivascular spaces and cerebral small vessel disease. *International journal of stroke* 10 (3), 376–381.
- Qi, H., Collins, S., Noble, A., 2017. Weakly supervised learning of placental ultrasound images with residual networks. In: *Proceedings of the Annual Conference on Medical Image Understanding and Analysis*. Springer, pp. 98–108.
- Redmon, J., Divvala, S., Girshick, R., Farhadi, A., 2016. You only look once: Unified, real-time object detection. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 779–788.
- Ren, M., Zemel, R.S., 2017. End-to-end instance segmentation with recurrent attention. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6656–6664.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-Net: Convolutional networks for biomedical image segmentation. In: *Proceedings of the International Conference on Medical image computing and computer-assisted intervention*. Springer, pp. 234–241.
- Sanginetto, E., Nabi, M., Culibrk, D., Sebe, N., 2018. Self paced deep learning for weakly supervised object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (1), 1–1.
- Schlegl, T., Seeböck, P., Waldstein, S.M., Schmidt-Erfurth, U., Langs, G., 2017. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In: *International Conference on Information Processing in Medical Imaging*. Springer, pp. 146–157.
- Schlemper, J., Oktay, O., Chen, L., Matthews, J., Knight, C., Kainz, B., Glocker, B., Rueckert, D., 2018. Attention-gated networks for improving ultrasound scan plane detection. *Medical Imaging with Deep Learning Conference*.
- Seguí, S., Pujol, O., Vitria, J., 2015. Learning to count with deep object features. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 90–96.
- Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D., 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In: *ICCV*, pp. 618–626.
- Sepehrband, F., Barisano, G., Sheikh-Bahaei, N., Cabeen, R.P., Choupan, J., Law, M., Toga, A.W., 2019. Image processing approaches to enhance perivascular space visibility and quantification using MRI. *Sci Rep* 9 (1), 1–12.
- Simonyan, K., Vedaldi, A., Zisserman, A., 2014. Deep inside convolutional networks: Visualising image classification models and saliency maps. In: *Proceedings of the International Conference for Learning Representations Workshop*.
- Springenberg, J.T., Dosovitskiy, A., Brox, T., Riedmiller, M., 2015. Striving for simplicity: The all convolutional net. In: *Proceedings of the International Conference for Learning Representations*.
- Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., Downey, P., Elliott, P., Green, J., Landray, M., et al., 2015. Uk biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS Med.* 12 (3), e1001779.
- Sudre, C.H., Anson, B.G., Ingala, S., Lane, C.D., Jimenez, D., Haider, L., Varsavsky, T., Smith, L., Jäger, H.R., Cardoso, M.J., 2018. 3D Multirater RCNN for multimodal multiclass detection and characterisation of extremely small objects. *arXiv preprint arXiv:1812.09046*.
- Tan, R.-x., Zhang, J., Chen, P., Wang, B., Xia, Y., 2018. Cells counting with convolutional neural network. In: *Proceedings of the International Conference on Intelligent Computing*. Springer, pp. 102–111.
- Walach, E., Wolf, L., 2016. Learning to count with cnn boosting. In: *Proceedings of the European conference on computer vision*. Springer, pp. 660–676.
- Wang, J., Knol, M., Tulpin, A., Dubost, F., De Bruijne, M., Vernooij, M., Adams, H., Ikram, M.A., Niessen, W., Roshchupkin, G., 2019. Grey matter age prediction as a biomarker for risk of dementia: a population-based study. *BioRxiv* 518506.
- Wardlaw, J.M., Smith, E.E., Biessels, G.J., Cordonnier, C., Fazekas, F., Frayne, R., Lindley, R.I., O'Brien, J., Barkhof, F., Benavente, O.R., et al., 2013. Neuroimaging standards for research into small vessel disease and its contribution to ageing and neurodegeneration. *The Lancet Neurology* 12 (8), 822–838.
- Xie, W., Noble, J.A., Zisserman, A., 2018. Microscopy cell counting and detection with fully convolutional regression networks. *Computer methods in biomechanics and biomedical engineering: Imaging & Visualization* 6 (3), 283–292.
- Xie, Y., Xing, F., Shi, X., Kong, X., Su, H., Yang, L., 2018. Efficient and robust cell detection: a structured regression approach. *Med Image Anal* 44, 245–254.
- Xue, Y., Ray, N., Hugh, J., Bigras, G., 2016. Cell counting by regression using convolutional neural network. In: *Proceedings of the European Conference on Computer Vision*. Springer, pp. 274–290.
- Zeiler, M.D., 2012. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.
- Zhang, Q., Wu, Y.N., Zhu, S.-C., 2018. Interpretable convolutional neural networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8827–8836.
- Zhang, Q.-s., Zhu, S.-C., 2018. Visual interpretability for deep learning: a survey. *Frontiers of Information Technology & Electronic Engineering* 19 (1), 27–39.
- Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A., 2016. Learning deep features for discriminative localization. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2921–2929.