

ToNNO: Tomographic Reconstruction of a Neural Network’s Output for Weakly Supervised Segmentation of 3D Medical Images

Marius Schmidt-Mengin¹ Alexis Benichoux¹ Shibeshih Belachew¹
 Nikos Komodakis^{1,3,4,5} Nikos Paragios^{1,2}

¹TheraPanacea, France ²CentraleSupélec, University of Paris-Saclay, France

³University of Crete, Greece ⁴IACM, FORTH, Greece ⁵Archimedes, Athena RC, Greece

Abstract

Annotating lots of 3D medical images for training segmentation models is time-consuming. The goal of weakly supervised semantic segmentation is to train segmentation models without using any ground truth segmentation masks. Our work addresses the case where only image-level categorical labels, indicating the presence or absence of a particular region of interest (such as tumours or lesions), are available. Most existing methods rely on class activation mapping (CAM). We propose a novel approach, **ToNNO**, which is based on the **Tomographic reconstruction of a Neural Network’s Output**. Our technique extracts stacks of slices with different angles from the input 3D volume, feeds these slices to a 2D encoder, and applies the inverse Radon transform in order to reconstruct a 3D heatmap of the encoder’s predictions. This generic method allows to perform dense prediction tasks on 3D volumes using any 2D image encoder. We apply it to weakly supervised medical image segmentation by training the 2D encoder to output high values for slices containing the regions of interest. We test it on four large scale medical image datasets and outperform 2D CAM methods. We then extend ToNNO by combining tomographic reconstruction with CAM methods, proposing *Averaged CAM* and *Tomographic CAM*, which obtain even better results.

1. Introduction

The advent of 3D medical imaging, such as computed tomography (CT), magnetic resonance imaging (MRI) or positron emission tomography (PET) has revolutionized clinical practice by enabling precise visualization and quantitative analysis of the internal structures of the body. Many scenarios involve segmenting particular regions of interest. For example, in radiotherapy workflows, it is necessary to segment organs and tumours in order to compute a treatment plan that maximizes the radiation dose




Figure 1. Our method, which allows to obtain high resolution segmentations without using any ground truth segmentation masks, is completely orthogonal to class activation mapping methods such as GradCAM [45] or LayerCAM [23].

received by the tumours while sparing the surrounding healthy tissues. In multiple sclerosis, segmentation of brain lesions allows to estimate lesion load for monitoring patients [40]. Furthermore, automatic segmentation methods may be useful as a verification step in order to detect tumours or lesions that were missed by human experts.

However, the volumetric nature of 3D medical images (also referred to as volumes) makes their manual segmentation slow and painstaking. Deep learning has become an important tool to automate this task. Deep learning techniques, though, typically still require large amounts of manually segmented training data, which can be prohibitive. Furthermore, manual segmentation can be subjective and vary a lot between different experts. Weakly supervised semantic segmentation methods aim to circumvent this issue by training semantic segmentation models without using any ground truth segmentation masks, relying instead on bounding boxes, scribbles, point-wise annotations, or image-level labels. Our work focuses on the latter.

To this end, we propose **ToNNO** (**Tomographic reconstruction of a Neural Network’s Output**), a generic method that produces a 3D heatmap which represents the predictions of a 2D encoder (such as a ResNet [16]) for

different slices of a given input volume. While many methods adopt an encoder-decoder structure, our method allows to perform dense prediction tasks on 3D images using only a 2D encoder. Furthermore, it enables to tap into the large amount of available 2D neural network implementations and pretrained weights. Our method is inspired by computed tomography. We extract 2D slices from a 3D image at many different angles, feed each of them to the 2D encoder, and use a tomographic reconstruction technique, namely the inverse Radon transform, in order to obtain a volumetric heatmap that represents the predicted slice logits in 3D.

We then apply this method to weakly supervised medical image semantic segmentation. Using only binary image-level labels that indicate whether or not the 3D images contain a given region of interest (e.g., tumour, lesion, etc.), we first train a 2D classifier to predict whether or not individual slices of these images contain the region of interest. Then, we apply ToNNO and show that it allows us to obtain a high resolution segmentation of the regions of interest, as shown in Figure 1.

ToNNO is orthogonal to class activation mapping (CAM) [64], which is currently the most common family of methods for weakly supervised medical image segmentation. As a baseline, we apply GradCAM [45] and LayerCAM [23] to the 2D classifier that we trained for ToNNO. We find that in most cases, ToNNO outperforms these two methods. Using the ideas behind ToNNO, we also propose to average the class activation maps produced by GradCAM [45] and LayerCAM [23] across many different angles, boosting the results of these methods by large amounts. Furthermore, we find that incorporating the filtering step—a key ingredient of the tomographic reconstruction technique that we use—into the averaging process allows to correct the inherent blurriness of class activation maps in order to obtain sharp averaged CAM heatmaps even for the deepest layers.

Thus, our contribution is threefold. First, we propose a framework to train a 2D classifier on slices extracted at any angle from the 3D volumes, using volume-level labels. Second, we propose ToNNO, a novel method to reconstruct 3D segmentations using this trained 2D classifier. Third, we incorporate CAM methods into our reconstruction framework, proposing Averaged CAM and Tomographic CAM.

2. Related Work

Weakly supervised semantic segmentation. This paradigm aims to segment 2D or 3D images without using ground truth segmentation masks. Multiple forms of supervision can be used instead: orthogonal segmentation masks [4], bounding boxes [8, 10, 24, 36, 42, 46, 50], scribbles [3, 28, 41, 51, 62, 63], point-wise annotations [2, 3, 35, 61], or, in its weakest and most challenging form,

image-level labels [1, 27, 38, 65]. When using image-level labels, most methods are based on deep learning and class activation mapping (CAM) [5, 23, 34, 45, 55, 64]. They make use of a classifier that was trained to classify between the different image-level labels and exploit activations and/or gradients of intermediate layers in order to obtain spatial information about the regions that cause the classifier to make its decision. For medical weakly supervised segmentation, CAM methods are most often applied to 2D slices [47, 52], sometimes deriving slice-level labels from 3D segmentation masks [6, 7, 14]. In this work, we apply GradCAM [45], and LayerCAM [23] which was shown to outperform GradCAM by large margins. One drawback of CAM methods is that the heatmaps that they generate usually have a resolution that is lower than the input.


3D medical image segmentation. It can be achieved with 3D encoder-decoder style neural networks [32], which maintain the 3D nature of the input data throughout the depth of the model. But 3D models can be hard to fit into the memory of GPUs for high resolution input data [20]. Thus, many of these methods rely on three 2D models that produce segmentations for each slice of the input volume in each viewing direction (axial, sagittal and coronal), which are then assembled back into a 3D segmentation [9, 20, 26, 43, 53, 59, 60]. In contrast, our method trains a single 2D encoder for slices which can have any 3D orientation.

Tomography. Tomographic reconstruction allows to recover volumetric information from slice-wise information. It is formalized by the Radon transform [18, 39] and most importantly, its inverse. Applications of tomographic reconstruction techniques include computed tomography (CT) imaging and electron tomography (ET). However, CT imaging is based on the 2D-1D Radon transform, meaning that 2D cross-sections of the patient are reconstructed using 1D integrals of the absorption of X-rays, and ET is based on either the 2D-1D or the 3D-1D Radon transform. These cases of the Radon transform have been extensively studied [13, 21, 48, 54]. Our work, however, uses the rather uncommon 3D-2D Radon transform, where a 3D volume is reconstructed using 2D slice integrals. To the best of our knowledge, no previous work has combined deep learning with the inverse Radon transform for dense prediction tasks on 3D images.


3. Method

3.1. Intuition

We assume that we have a dataset of volumes with associated binary labels $\{(V_1, y_1), \dots, (V_N, y_n)\}$, where $\forall i, V_i \in \mathbb{R}^{d_i \times h_i \times w_i}$ and $y_i \in \{0, 1\}$. Our goal is to obtain a



(a) Histograms of logits for 100,000 slices of the MosMed dataset (validation set) extracted from multiple positive and negative volumes.



(b) Logits produced by the trained classifier for each slice of the volume shown in the background (Duke dataset) along the vertical and horizontal axis allow to pinpoint the location of the tumour.

Figure 2. Visualisation of logits produced by the 2D classifier g_θ

segmentation of the regions that distinguish the positive class from the negative class. For example, if the positive class consists of images of cancer patients and the negative class consists of images of healthy patients, our goal is to segment the tumours. Our method starts by training a 2D classification neural network g_θ , with parameters θ , to distinguish between slices of positive volumes and slices of negative volumes. In order to do that, we sample random slices (which can have any 3D orientation) from the training volumes and associate to each slice the ground truth label of the volume from which it came. Note that this effectively introduces label noise, as in positive volumes, all slices get assigned a positive label even if they don't contain any regions of interest. But we have no way of deriving more precise labels, as per the problem definition, only volume-level labels are available. We then optimize θ to minimize the binary cross-entropy loss of each slice with respect to its assigned label.

After training, the distributions of the logits produced by the classifier for slices coming from the positive and negative classes are different, as shown in Figure 2a. When feeding each slice of a given volume to the classifier, we obtain a logit profile along the slice axis that can allow us to pinpoint the location of the regions of interest along this axis (see Figure 2b). If we repeat this process for different slice axes, we can narrow down the location in three dimensions. This process is formalized by the theory of tomographic reconstruction, whose principal mathematical tool is the Radon transform.

Our weakly supervised segmentation method thus consists in two steps. First, we train the classifier with stacks of slices that can have any 3D orientation, which can be

seen as applying random 3D rotations to the volumes during training. Then, we reconstruct a heatmap by applying the inverse Radon transform to logits obtained for stacks of slices across many different orientations.

3.2. Radon transform and its inverse

Let $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ be a 3D function. For the sake of simplicity, we assume it to be infinitely differentiable with compact support. Let \mathbb{S}^2 be the unit sphere of \mathbb{R}^3 . For a unit vector $\hat{\mathbf{n}} \in \mathbb{S}^2$ and $s \in \mathbb{R}$, let $\xi(s, \hat{\mathbf{n}}) = \{\mathbf{x} \in \mathbb{R}^3, \mathbf{x} \cdot \hat{\mathbf{n}} = s\}$ be the plane orthogonal to $\hat{\mathbf{n}}$ containing the point $s\hat{\mathbf{n}}$. We define the 3D-2D Radon transform of f as the function $\mathcal{R}f : \mathbb{R} \times \mathbb{S}^2 \rightarrow \mathbb{R}$ by

$$\mathcal{R}f(s, \hat{\mathbf{n}}) := \int_{\mathbf{x} \in \xi(s, \hat{\mathbf{n}})} f(\mathbf{x}) \, d\mathbf{x} \quad (1)$$

It can be shown [18, Theorem 3.6] that f can be recovered from $\mathcal{R}f$ using the filtered backprojection formula:

$$f(\mathbf{x}) = c \int_{\hat{\mathbf{n}} \in \mathbb{S}^2} \frac{\partial^2 \mathcal{R}f}{\partial s^2}(\mathbf{x} \cdot \hat{\mathbf{n}}, \hat{\mathbf{n}}) \, d\hat{\mathbf{n}} \quad (2)$$

where c is a negative constant. For a function $\varphi : \mathbb{R} \times \mathbb{S}^2 \rightarrow \mathbb{R}$, we thus define $\mathcal{R}^{-1}\varphi : \mathbb{R}^3 \rightarrow \mathbb{R}$ by

$$\mathcal{R}^{-1}\varphi(\mathbf{x}) := c \int_{\hat{\mathbf{n}} \in \mathbb{S}^2} \frac{\partial^2 \varphi}{\partial s^2}(\mathbf{x} \cdot \hat{\mathbf{n}}, \hat{\mathbf{n}}) \, d\hat{\mathbf{n}} \quad (3)$$

It holds that for all $\hat{\mathbf{n}}, \mathbf{x} \in \xi(\mathbf{x} \cdot \hat{\mathbf{n}}, \hat{\mathbf{n}})$. Thus, for a given \mathbf{x} , this integral integrates the function $\partial^2 \varphi / \partial s^2$ over all planes containing the point \mathbf{x} and is called the backprojection operator.

3.3. Implementation

Now, let $V : [d] \times [h] \times [w] \rightarrow \mathbb{R}$ (where for an integer n , we define $[n] = \{0, 1, \dots, n-1\}$) be a discrete 3D volume, such as a medical image. The Radon transform cannot be directly applied to V , as V is not defined on \mathbb{R}^3 . Let $\tilde{V} : \mathbb{R}^3 \rightarrow \mathbb{R}$ be an interpolation of V , such that $\tilde{V}(-1, -1, -1) = V(0, 0, 0)$ and $\tilde{V}(1, 1, 1) = V(d-1, h-1, w-1)$, and $\tilde{V}(\mathbf{x}) = 0$ if $\mathbf{x} \notin [-1, 1]^3$. The Radon transform of \tilde{V} is well-defined, but not straightforward to compute numerically. We can approximate $\mathcal{R}\tilde{V}(s, \hat{\mathbf{n}})$ by extracting a discrete 2D slice of \tilde{V} and summing its pixels. For that, let \mathbf{u} and \mathbf{v} be two vectors orthogonal to $\hat{\mathbf{n}}$. We define the slice $S_V(s, \hat{\mathbf{n}}, \mathbf{u}, \mathbf{v}) : [h_S] \times [w_S] \rightarrow \mathbb{R}$ by

$$S_V(s, \hat{\mathbf{n}}, \mathbf{u}, \mathbf{v})(i, j) := \tilde{V}(s\hat{\mathbf{n}} + i'\mathbf{u} + j'\mathbf{v}) \quad (4)$$

where $i' := 2\frac{i}{h_S-1} - 1$ and $j' := 2\frac{j}{w_S-1} - 1$ map integer coordinates i and j to the range $[-1, 1]$. $S_V(s, \hat{\mathbf{n}}, \mathbf{u}, \mathbf{v})$ is a 2D image, which can be efficiently extracted with the function `grid_sample` of PyTorch [37] to be used as input for a neural network. It is the slice of V which is oriented by

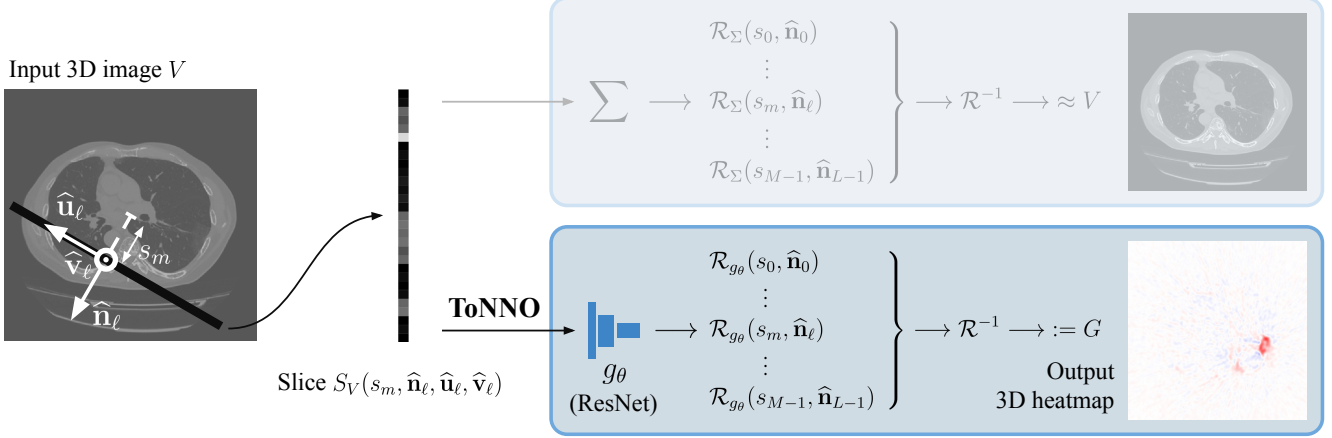


Figure 3. **Overview of ToNNO.** 3D volumes are represented as 2D images and 2D slices as 1D rows of pixels. First, slices are extracted from the input volume (left). By summing the pixels of each slice (top), we can approximate the Radon transform of the input volume V , and applying the inverse Radon transform \mathcal{R}^{-1} allows to reconstruct V . The idea of this work is to replace the sum by a trained neural network before applying the inverse Radon transform (bottom).

\mathbf{u} and \mathbf{v} , whose normal vector is $\hat{\mathbf{n}}$, and which contains the point $s\hat{\mathbf{n}}$. The lengths of \mathbf{u} and \mathbf{v} determine the scale of the slice and the angle between them controls the shear (which allows us to conveniently introduce augmentations during training). For tomographic reconstruction, we choose them perpendicular to each other, with unit length. Therefore, let $\hat{\mathbf{n}} \mapsto (\hat{\mathbf{u}}(\hat{\mathbf{n}}), \hat{\mathbf{v}}(\hat{\mathbf{n}}))$ be a function such that $(\hat{\mathbf{n}}, \hat{\mathbf{u}}(\hat{\mathbf{n}}), \hat{\mathbf{v}}(\hat{\mathbf{n}}))$ is an orthonormal basis (see Appendix A.1). We define the approximate Radon transform of V as

$$\mathcal{R}_\Sigma V(s, \hat{\mathbf{n}}) := \frac{1}{h_S w_S} \sum_{i=0}^{h_S-1} \sum_{j=0}^{w_S-1} S_V(s, \hat{\mathbf{n}}, \hat{\mathbf{u}}(\hat{\mathbf{n}}), \hat{\mathbf{v}}(\hat{\mathbf{n}}))(i, j) \approx \mathcal{R}\tilde{V}(s, \hat{\mathbf{n}}) \quad (5)$$

3.4. ToNNO

The idea of this work is to replace the sum of Equation (5) with a trained 2D neural network that maps a 2D image to a real number, before applying the inverse Radon transform (see Figure 3). We define

$$\mathcal{R}_{g_\theta} V(s, \hat{\mathbf{n}}) := g_\theta(S_V(s, \hat{\mathbf{n}}, \hat{\mathbf{u}}(\hat{\mathbf{n}}), \hat{\mathbf{v}}(\hat{\mathbf{n}}))) \quad (6)$$

The output of our method is a volume $G : [D] \times [H] \times [W] \rightarrow \mathbb{R}$ (in practice, to facilitate evaluation, we set the output shape to be equal to the input shape, that is, $D = d$, $H = h$ and $W = w$) such that

$$G(i, j, k) \approx \mathcal{R}^{-1}(\mathcal{R}_{g_\theta} V)(i', j', k') \quad (7)$$

where $i' = 2\frac{i}{D-1} - 1$, $j' = 2\frac{j}{H-1} - 1$ and $k' = 2\frac{k}{W-1} - 1$ map integer coordinates to the range $[-1, 1]$.

The integral over $\hat{\mathbf{n}} \in \mathbb{S}^2$ of Equation (3) can be approximated by a sum over a set of L unit vectors that

are evenly distributed over the unit sphere \mathbb{S}^2 . To get such a set of unit vectors, we use the Fibonacci lattice: for $\ell \in [L]$, let $y_\ell = 2\frac{\ell}{L-1} - 1 \in [-1, 1]$, $r_\ell = \sqrt{1 - y_\ell^2}$, $\theta_\ell = \ell\alpha$ where $\alpha = \pi(3 - \sqrt{5})$ is the golden angle, $x_\ell = r_\ell \cos \theta_\ell$ and $z_\ell = r_\ell \sin \theta_\ell$. We then set $\hat{\mathbf{n}}_\ell = (x_\ell, y_\ell, z_\ell)$. We then have

$$\mathcal{R}^{-1}(\mathcal{R}_{g_\theta} V)(i', j', k') \approx \frac{c|\mathbb{S}^2|}{L} \sum_{\ell=0}^{L-1} \frac{\partial^2 \mathcal{R}_{g_\theta} V}{\partial s^2} \left(\begin{pmatrix} i' \\ j' \\ k' \end{pmatrix} \cdot \hat{\mathbf{n}}_\ell, \hat{\mathbf{n}}_\ell \right) \quad (8)$$

where $|\mathbb{S}^2|$ is the surface area of the unit sphere.

It remains to estimate the quantity $\partial^2 \mathcal{R}_{g_\theta} V(s, \hat{\mathbf{n}}_\ell) / \partial s^2$. In order to do that, for a fixed unit vector $\hat{\mathbf{n}}_\ell$, we sample the function $s \mapsto \mathcal{R}_{g_\theta} V(s, \hat{\mathbf{n}}_\ell)$ at $M + 2$ regular intervals $s_m = 2\frac{m}{M-1} - 1$, $m \in \{-1, 0, \dots, M\}$, which corresponds to extracting a stack of $M + 2$ regularly spaced slices with normal vector $\hat{\mathbf{n}}_\ell$ and feeding each of them to the neural network. Letting $p_m(\hat{\mathbf{n}}_\ell) := \mathcal{R}_{g_\theta} V(s_m, \hat{\mathbf{n}}_\ell)$, we use finite differences to estimate the second derivative, defining, for $m \in \{0, \dots, M-1\}$

$$q_m(\hat{\mathbf{n}}_\ell) := \frac{p_{m+1}(\hat{\mathbf{n}}_\ell) + p_{m-1}(\hat{\mathbf{n}}_\ell) - 2p_m(\hat{\mathbf{n}}_\ell)}{1/(M-1)} \quad (9)$$

so that

$$q_m(\hat{\mathbf{n}}_\ell) \approx \frac{\partial^2 \mathcal{R}_{g_\theta} V}{\partial s^2}(s_m, \hat{\mathbf{n}}_\ell) \quad (10)$$

Now, let $s \mapsto \tilde{q}(s, \hat{\mathbf{n}}_\ell)$ be an interpolation of $m \mapsto q_m(\hat{\mathbf{n}}_\ell)$, such that $\tilde{q}(-1, \hat{\mathbf{n}}_\ell) = q_0(\hat{\mathbf{n}}_\ell)$, $\tilde{q}(1, \hat{\mathbf{n}}_\ell) = q_{M-1}(\hat{\mathbf{n}}_\ell)$ and $\tilde{q}(s, \hat{\mathbf{n}}_\ell) = 0$ for $s \notin [-1, 1]$ (which again, can be efficiently computed using the function `grid_sample` of PyTorch).




Figure 4. Input reconstructions for different average pooling kernel sizes.

We finally define

$$G(i, j, k) := \frac{c|\mathbb{S}^2|}{L} \sum_{l=0}^{L-1} \tilde{q} \left(\begin{pmatrix} i' \\ j' \\ k' \end{pmatrix} \cdot \hat{\mathbf{n}}_\ell, \hat{\mathbf{n}}_\ell \right) \quad (11)$$

$$\approx \mathcal{R}^{-1}(\mathcal{R}_{g_\theta} V)(i', j', k')$$

The correctness of the implementation can be verified by substituting \mathcal{R}_{g_θ} for \mathcal{R}_Σ (Equation (5)) and checking that in this case, $G \approx V$.

3.5. Averaged CAM and Tomographic CAM

Class activation mapping (CAM) methods such as GradCAM [45] and LayerCAM [23] are currently the most common method for weakly supervised semantic segmentation. For a given slice, they provide a heatmap of shape $h_{\text{CAM}} \times w_{\text{CAM}}$, examples of which are shown in columns 3 and 5 of Figure 6. h_{CAM} and w_{CAM} correspond to the resolution of the features at the layer from which the class activation maps are extracted. For a ResNet with input shape 224×224 , the resolution at the i -th layer is $h_{\text{CAM}} = w_{\text{CAM}} = 224/2^{i+1}$. We propose to average these CAM heatmaps across multiple rotations. For a given normal vector $\hat{\mathbf{n}}_\ell$, we extract the $M+2$ slices $S_V(s_m, \hat{\mathbf{n}}_\ell, \hat{\mathbf{u}}_\ell, \hat{\mathbf{v}}_\ell)$ and obtain the heatmap for each slice by applying LayerCAM to the neural network g_θ . We denote by $p_m^{ab}(\hat{\mathbf{n}}_\ell)$ the value of the location $(a, b) \in [h_{\text{CAM}}] \times [w_{\text{CAM}}]$ of the m -th heatmap. Let $(s, u, v) \mapsto \tilde{p}(s, u, v, \hat{\mathbf{n}}_\ell)$ be an interpolation of the 3D volume $(m, a, b) \mapsto p_m^{ab}(\hat{\mathbf{n}}_\ell)$. A point with coordinates $(i' j' k')$ in the canonical coordinate system has coordinates $s = (i' j' k') \cdot \hat{\mathbf{n}}_\ell$, $u = (i' j' k') \cdot \hat{\mathbf{u}}_\ell$, $v = (i' j' k') \cdot \hat{\mathbf{v}}_\ell$ in the coordinate system $(\hat{\mathbf{n}}_\ell, \hat{\mathbf{u}}_\ell, \hat{\mathbf{v}}_\ell)$. We thus define the average heatmap, that we call **Averaged CAM**:

$$G_{\text{avgCAM}}(i, j, k) \quad (12)$$

$$:= \frac{c|\mathbb{S}^2|}{L} \sum_{l=0}^{L-1} \tilde{p} \left(\begin{pmatrix} i' \\ j' \\ k' \end{pmatrix} \cdot \hat{\mathbf{n}}_\ell, \begin{pmatrix} i' \\ j' \end{pmatrix} \cdot \hat{\mathbf{u}}_\ell, \begin{pmatrix} i' \\ j' \\ k' \end{pmatrix} \cdot \hat{\mathbf{v}}_\ell, \hat{\mathbf{n}}_\ell \right)$$

Using class activation maps from deeper layers (and thus of lower spatial resolution) leads to blurrier results, as can be seen in Figures 11, 13, 15, and 17 in the Appendix. Figure 5 confirms that the performance of Averaged LayerCAM quickly deteriorates when it is applied at deeper layers. We find that we can partially correct this by applying the second derivative along the slice axis separately for each location (a, b) , that is, we define q_m^{ab} by replacing p_m by p_m^{ab} in Equation (9). We then define **Tomographic CAM**, by substituting \tilde{p} with \tilde{q} (which is the interpolation of q , as defined for p) in Equation (12):

$$G'_{\text{avgCAM}}(i, j, k) \quad (13)$$

$$:= \frac{c|\mathbb{S}^2|}{L} \sum_{l=0}^{L-1} \tilde{q} \left(\begin{pmatrix} i' \\ j' \\ k' \end{pmatrix} \cdot \hat{\mathbf{n}}_\ell, \begin{pmatrix} i' \\ j' \end{pmatrix} \cdot \hat{\mathbf{u}}_\ell, \begin{pmatrix} i' \\ j' \\ k' \end{pmatrix} \cdot \hat{\mathbf{v}}_\ell, \hat{\mathbf{n}}_\ell \right)$$

Figure 5 shows that Tomographic LayerCAM indeed allows to maintain good performance for deeper layers.

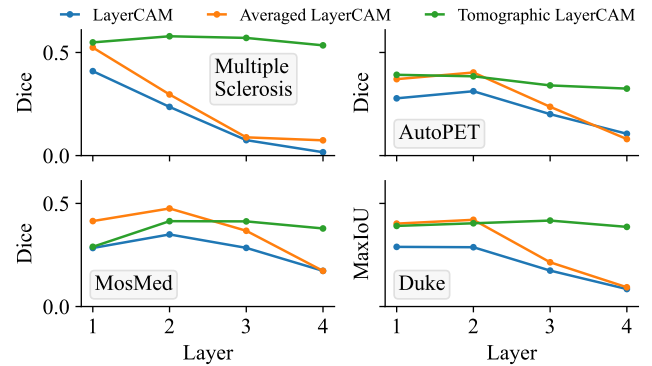


Figure 5. Dice score for different layers of LayerCAM, Averaged LayerCAM and Tomographic LayerCAM.

For gaining further insight into this, we can assume that the value at position $(a, b) \in [h_{\text{CAM}}] \times [w_{\text{CAM}}]$ roughly represents the amount of activation for a region of shape $224/h_{\text{CAM}} \times 224/w_{\text{CAM}}$ centered on the pixel

$(a \cdot 224/h_{\text{CAM}}, b \cdot 224/w_{\text{CAM}})$ in the input image. Thus, we can obtain a "CAM" version of \mathcal{R}_Σ of Equation (5) by average-pooling the input image with a kernel size and stride of $(224/h_{\text{CAM}}, 224/w_{\text{CAM}})$. We can then use the formulas of Averaged CAM and Tomographic CAM (Equations (12) and (13)) to reconstruct the input image. The results are presented on Figure 4. As can be seen, when the CAM maps have high resolution, then simple averaging does allow to recover the input. But when the CAM maps have lower resolution, incorporating the second derivative becomes necessary. In between, the "Averaged" reconstructions are blurry, and the "Tomographic" reconstructions (i.e., with second derivative) are too sharp, as if a filter had been applied on the image. Future work could look into a way of reconstructing these intermediate cases, for example using iterative reconstruction techniques.

3.6. Hyperparameters

Tomographic reconstruction We experiment with different numbers of slices (M in Equation (9)) and angles (L in Equation (8)). Reconstructions of one sample of the MosMed dataset are shown in Figure 8 of the Appendix for $(M, L) \in \{25, 50, 100, 200\} \times \{500, 1000, 2000, 5000\}$. Increasing the number of slices M increases the sharpness of the reconstruction but also increases the amount of noise. The number of angles L reduces the noise. In our experiments, we choose $M = 100$ and $L = 2000$ and we set the output shape to be equal to the shape of the input volume. For Averaged CAM and Tomographic CAM, we set $L = 1000$ because obtaining the class activation maps involves more computations. All interpolations are performed with trilinear interpolation. With these parameters, the reconstruction time is around 40 seconds for ToNNO and 1 minute for Averaged and Tomographic CAM on a NVIDIA GeForce RTX 3090 GPU. This is slower than slice-wise CAM methods, which take less than a second per volume on a GPU.

Classifier choice The neural network should have three characteristics: fast inference (as many inferences are needed for tomographic reconstruction), availability of pre-trained weights, and absence of batch normalization (BN) layers [22]. We hypothesize that the latter is important because BN breaks the independence between samples within a batch, leading to a low-loss but less informative solution, as reported in [17, 19]. For example, if one slice is identified by the classifier as containing a discriminative region, then intra-batch information leakage caused by BN could enable the classifier to also output high logits for slices of the volume that do not contain any discriminative region, leading to a lower loss but useless predictions. In the end, we select two ResNets [16, 56] from the `timm` [57] library, for which ImageNet [11] pretrained weights are available:

- The **ResNet-10-T** [56] which is a lightweight ResNet. It comes with batch normalization layers, so we explore two modifications: freezing them or replacing them with group normalization layers [58] (still initializing the other layers with the pretrained weights).
- The **ResNet-50** [16] with group normalization. Because of heavier overfitting, we train it for only 40% of the number of iterations of the ResNet-10-T.

We experimented with more recent backbones, such as EfficientNet [49], ConvNeXt [30], ViT [12] and Swin [29], but were unable to surpass the more simple ResNets. Careful tuning of training hyper-parameters may allow them to achieve better performance but this is not the focus of our work.

Our results in Table 1 show that the ResNet-10-T with group normalization and frozen batch normalization are on par, and better than with vanilla batch normalization. We did not find the ResNet-50 to be significantly better than the ResNet-10-T. The ResNet-10-T with random initialization achieves much lower performance, confirming the benefit of using pretrained weights. In all the other experiments presented in this paper, we use the ResNet-10-T with frozen batch normalization and ImageNet pretrained weights.

Model	F1-score	Dice/IoU	BA
R-10-T, batch norm	0.48	0.31	0.84
R-10-T, frozen batch norm	0.55	0.39	0.84
R-10-T, group norm	0.54	0.36	0.85
R-10-T, not pretrained	0.28	0.27	0.68
R-50, group norm	0.56	0.40	0.84

Table 1. **Comparison of different model configurations.** The metrics were averaged over the four datasets. Per-dataset results are presented in Appendix A.5.

4. Experiments

4.1. Datasets

We use 4 different datasets for our experiments:

- **Multiple Sclerosis (MS) (brain MRI)** This large private dataset consists of 9113 studies. The goal is to segment gadolinium-enhancing lesions, which are typically hyperintense on the post-gadolinium injection T1 weighted scan. This abnormal gadolinium uptake is due to blood-brain barrier breakdown and is a sign of acute multiple sclerosis activity.
- **AutoPET-II (whole body PET-CT)** This public dataset [15] consists of 1,014 FDG-PET/CT pairs. The objective of this exam is to segment FDG-avid tumours.
- **MosMedData COVID-19 (thoracic CT)** This public dataset [33] consists of 1110 studies of thoracic CT scans. The goal is to segment pulmonary COVID-19 lesions.

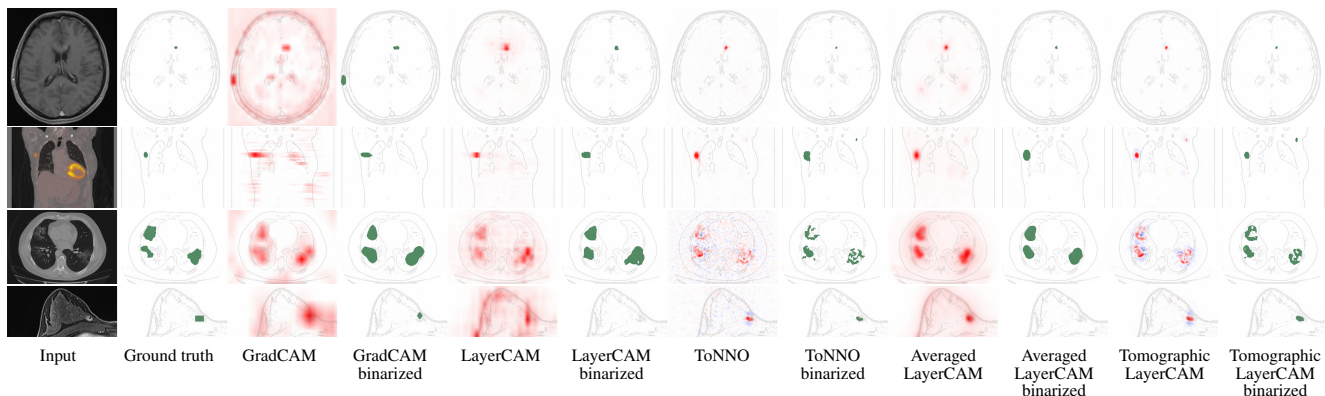


Figure 6. **Examples of heatmaps and segmentations obtained with each method.** From top to bottom: Multiple Sclerosis, AutoPET, MosMed and Duke datasets. The images were cropped. For GradCAM, LayerCAM, Averaged LayerCAM and Tomographic LayerCAM, we chose the layer which gave the best results. The binarized heatmaps are obtained using the process described in the Appendix A.7. In the heatmaps, blue represents negative values and red represents positive values. More examples are provided in the Appendix.

- **Duke breast cancer MRI** This public dataset [44] consists of breast MRIs of 922 biopsy-confirmed invasive breast cancer patients. The goal is to detect the tumours.

We split each dataset into training and validation sets at the patient level. Dataset characteristics are presented in Appendix A.3 and preprocessing details are explained in Appendix A.4. For the Multiple Sclerosis and AutoPET datasets, ground truth segmentation masks are available for all samples. Patients with non-empty segmentation mask are considered positive while the rest are considered negative. For the MosMed dataset, ground truth segmentation masks are available for 50 patients, which we all put in the validation set. The positive/negative labels of all patient are available separately. For the Duke dataset, all patients are positive for breast cancer, and exactly one bounding box per patient, delimiting the main tumour, is available. Our method needs both positive and negative samples, so we separate the left and right breast of each patient and consider the breast with bounding box to be positive.

4.2. Training procedure

During a training step, we sample a batch of B samples $\{(V_{i_1}, y_{i_1}), \dots, (V_{i_B}, y_{i_B})\}$ with replacement from our training dataset. We then extract M_{train} randomly oriented slices of shape 224×224 from each volume and apply augmentations. We concatenate the $B \times M_{\text{train}}$ slices in a batch and perform a training step by associating to each slice the label of the volume from which it came. More details can be found in Appendix A.6.

4.3. Baselines

We compare our method to class activation methods, namely GradCAM and LayerCAM. We apply them to the ResNet-10-T with frozen batch normalization at the output

of layers 1, 2, 3 and 4, to each axial slice of the input volume, and obtain the output volume by stacking the class activation maps back along the axial axis. To be input to the classifier, each slice is resampled to match the input shape of the neural network (224×224), and the resulting class activation map is upsampled to match the original slice shape. The output volume thus has the same shape as the input volume.

In Appendix A.8, we also provide results for GradCAM++ [5] and ScoreCAM [55]. Furthermore, we tried to apply the CAM methods along the 3 spatial axes, averaging the outputs. In some cases, this resulted in significant improvements (Appendix A.8).

4.4. Evaluation

In order to compute segmentation metrics, we first binarize the heatmaps using the process described in Appendix A.7. We then compute the connected components of the binarized heatmaps and the ground truth segmentation and report precision, recall and F1-score. We also report the dice score. For the Duke dataset, as no ground truth segmentation masks are available (we only have one bounding box annotation per patient), we replace the dice score with the intersection-over-union of the predicted bounding box which best matches the ground truth bounding box. In the different tables, this metric is called MaxIoU. As there is at most one ground truth bounding box per patient, despite there sometimes being multiple tumours, many predicted connected components that are considered as false positives may actually be true positives. Thus, precision and F1-score are less relevant for this dataset.

4.5. Results

Main results are summarized in Table 2, and results for each method applied at each layer can be found in the Appendix. ToNNO outperforms LayerCAM (which largely outperforms

Dataset	Method	Precision	Recall	F1-score	Dice/MaxIoU	Balanced accuracy
Multiple Sclerosis	LayerCAM (layer 1)	0.76	0.78	0.73	0.41	0.89
	ToNNO	0.82 (+0.06)	0.88 (+0.10)	0.82 (+0.09)	0.53 (+0.12)	0.90 (+0.01)
	Averaged LayerCAM (layer 1)	0.86 (+0.10)	0.87 (+0.09)	0.84 (+0.11)	0.52 (+0.11)	0.91 (+0.02)
	Tomographic LayerCAM (layer 2)	0.87 (+0.11)	0.87 (+0.09)	0.84 (+0.11)	0.58 (+0.17)	0.94 (+0.05)
AutoPET	LayerCAM (layer 2)	0.59	0.29	0.35	0.31	0.80
	ToNNO	0.52 (-0.07)	0.41 (+0.12)	0.39 (+0.04)	0.32 (+0.01)	0.74 (-0.06)
	Averaged LayerCAM (layer 2)	0.65 (+0.06)	0.33 (+0.04)	0.40 (+0.05)	0.40 (+0.09)	0.83 (+0.03)
	Tomographic LayerCAM (layer 1)	0.68 (+0.09)	0.47 (+0.18)	0.49 (+0.14)	0.39 (+0.08)	0.74 (-0.06)
MosMed	LayerCAM (layer 2)	0.68	0.32	0.39	0.35	0.90
	ToNNO	0.69 (+0.01)	0.45 (+0.13)	0.50 (+0.11)	0.35	0.93 (+0.03)
	Averaged LayerCAM (layer 2)	0.72 (+0.04)	0.49 (+0.17)	0.53 (+0.14)	0.48 (+0.13)	0.95 (+0.05)
	Tomographic LayerCAM (layer 3)	0.72 (+0.04)	0.52 (+0.20)	0.55 (+0.16)	0.41 (+0.06)	0.89 (-0.01)
Duke	LayerCAM (layer 2)	0.17	0.81	0.24	0.29	0.58
	ToNNO	0.43 (+0.26)	0.77 (-0.04)	0.51 (+0.27)	0.37 (+0.08)	0.79 (+0.21)
	Averaged LayerCAM (layer 2)	0.37 (+0.20)	0.84 (+0.03)	0.47 (+0.23)	0.42 (+0.13)	0.75 (+0.17)
	Tomographic LayerCAM (layer 3)	0.41 (+0.24)	0.81 (=)	0.51 (+0.27)	0.42 (+0.13)	0.79 (+0.21)

Table 2. **Main quantitative results of our method and best performing baseline method (LayerCAM).** For LayerCAM, Averaged LayerCAM and Tomographic LayerCAM, we selected the layer which provided the best results. Detailed results for each layer are provided in the Appendix. The numbers in color indicate the difference with LayerCAM.

GradCAM as can be seen in the Appendix) in most cases. Averaged LayerCAM and Tomographic LayerCAM achieve even better results. Tomographic LayerCAM achieves the best results in terms of F1-score, with an average improvement of 0.14 over LayerCAM. In terms of recall, Tomographic LayerCAM achieves particularly large gains over LayerCAM for the AutoPET and MosMed datasets.

All methods achieve their best results on the Multiple Sclerosis dataset. This could be attributed to several factors: the training dataset is an order of magnitude larger than the other datasets and the shape of the lesions is simple. Our method can be compared to the results presented in [35], which uses point-wise lesion annotations while we use only image-level labels. Although they do not use the same dataset and the same evaluation process, the results that we achieve are within the same order of magnitude as theirs (F1-score of 0.86 for their method versus 0.84 for Tomographic LayerCAM).

Examples of heatmaps and associated binary segmentations for each method and dataset are shown in Figure 6, and examples of heatmaps for each layer can be found in the Appendix. ToNNO and Tomographic LayerCAM provide much sharper heatmaps than baseline 2D CAM methods and Averaged CAM. While this is an advantage for segmenting the small lesions of the Multiple Sclerosis dataset, it results in undersegmentation of the large lesions of the MosMed dataset.

5. Conclusion and future work

In this work, we have presented ToNNO, a novel method that allows to perform dense prediction tasks on 3D images using a 2D encoder. We have shown that in the case where the encoder is trained on a weakly supervised classification task, our method can achieve better results than state-of-the-art CAM methods. We have also proposed Averaged and Tomographic CAM, which allow to integrate CAM methods into our framework in order to achieve even better results.

We believe that ToNNO could be applied in other ways. For example, it could allow to obtain dense features for 3D images using an encoder that is pretrained either on large scale computer vision datasets such as ImageNet [11], or by using self-supervised learning techniques directly on slices of 3D medical image datasets. ToNNO, contrary to CAM methods, could also be used to obtain quantitative heatmaps in the case where the encoder is pretrained in a quantitative way, for example on a regression task.

Future work could concentrate on finding a canonical way—for example using iterative methods—to perform tomographic reconstruction with class activation maps, as we have seen that neither Tomographic CAM nor Averaged CAM are perfect generalizations of ToNNO.

References

- [1] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *CVPR*, pages 4981–4990. Computer Vision Foundation / IEEE Computer Society, 2018.

- [2] Amy L. Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What's the point: Semantic segmentation with point supervision. In *ECCV (7)*, pages 549–565. Springer, 2016. [2](#)
- [3] Yuri Boykov and Marie-Pierre Jolly. Interactive graph cuts for optimal boundary and region segmentation of objects in N-D images. In *ICCV*, pages 105–112. IEEE Computer Society, 2001. [2](#)
- [4] Heng Cai, Shumeng Li, Lei Qi, Qian Yu, Yinghuan Shi, and Yang Gao. Orthogonal annotation benefits barely-supervised medical image segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 3302–3311, 2023. [2](#)
- [5] Aditya Chattopadhyay, Anirban Sarkar, Prantik Howlader, and Vineeth N. Balasubramanian. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In *WACV*, pages 839–847. IEEE Computer Society, 2018. [2](#), [7](#), [14](#)
- [6] Yu-Jen Chen, Xinrong Hu, Yiyu Shi, and Tsung-Yi Ho. AME-CAM: attentive multiple-exit CAM for weakly supervised segmentation on MRI brain tumor. In *MICCAI (1)*, pages 173–182. Springer, 2023. [2](#)
- [7] Zhang Chen, Zhiqiang Tian, Jihua Zhu, Ce Li, and Shaoyi Du. C-CAM: causal CAM for weakly supervised semantic segmentation on medical image. In *CVPR*, pages 11666–11675. IEEE, 2022. [2](#)
- [8] Julian Chibane, Francis Engelmann, Tuan Anh Tran, and Gerard Pons-Moll. Box2mask: Weakly supervised 3d semantic instance segmentation using bounding boxes. In *ECCV (31)*, pages 681–699. Springer, 2022. [2](#)
- [9] Hejie Cui, Xinglong Liu, and Ning Huang. Pulmonary vessel segmentation based on orthogonal fused u-net++ of chest CT images. In *Medical Image Computing and Computer Assisted Intervention - MICCAI 2019 - 22nd International Conference, Shenzhen, China, October 13-17, 2019, Proceedings, Part VI*, pages 293–300. Springer, 2019. [2](#)
- [10] Jifeng Dai, Kaiming He, and Jian Sun. Boxesup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 1635–1643, 2015. [2](#)
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. [6](#), [8](#)
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. [6](#)
- [13] J. Frank. *Electron Tomography: Methods for Three-Dimensional Visualization of Structures in the Cell*. Springer New York, 2008. [2](#)
- [14] Marcel Früh, Marc Fischer, Andreas Schilling, Sergios Gatidis, and Tobias Hepp. Weakly supervised segmentation of tumor lesions in PET-CT hybrid imaging. *J. Med. Imaging (Bellingham)*, 8(5):054003, 2021. [2](#)
- [15] Sergios Gatidis and Thomas Kuestner. A whole-body FDG-PET/CT dataset with manually annotated tumor lesions, 2022. [6](#), [12](#)
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [1](#), [6](#)
- [17] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. [6](#)
- [18] Sigurdur Helgason. *The Radon Transform*. Springer US, 1999. [2](#), [3](#)
- [19] Olivier J. Hénaff. Data-efficient image recognition with contrastive predictive coding. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, pages 4182–4192. PMLR, 2020. [6](#)
- [20] Leonie Henschel, Sailesh Conjeti, Santiago Estrada, Kersten Diers, Bruce Fischl, and Martin Reuter. Fastsurfer - A fast and accurate deep learning based neuroimaging pipeline. *NeuroImage*, 219:117012, 2020. [2](#)
- [21] G.T. Herman. *Fundamentals of Computerized Tomography: Image Reconstruction from Projections*. Springer London, 2009. [2](#)
- [22] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015. [6](#)
- [23] Peng-Tao Jiang, Chang-Bin Zhang, Qibin Hou, Ming-Ming Cheng, and Yunchao Wei. Layercam: Exploring hierarchical class activation maps for localization. *IEEE Transactions on Image Processing*, 30:5875–5888, 2021. [1](#), [2](#), [5](#)
- [24] Anna Khoreva, Rodrigo Benenson, Jan Hendrik Hosang, Matthias Hein, and Bernt Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *CVPR*, pages 1665–1674. IEEE Computer Society, 2017. [2](#)
- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [14](#)
- [26] Titinunt Kitrungratsakul, Xian-Hua Han, Yutarō Iwamoto, Lanfen Lin, Amir Hossein Foruzan, Wei Xiong, and Yen-Wei Chen. VesselNet: A deep convolutional neural network with multi pathways for robust hepatic vessel segmentation. *Computerized Medical Imaging and Graphics*, 75:74–83, 2019. [2](#)
- [27] Alexander Kolesnikov and Christoph H. Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *ECCV (4)*, pages 695–711. Springer, 2016. [2](#)
- [28] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *CVPR*, pages 3159–3167. IEEE Computer Society, 2016. [2](#)

- [29] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 9992–10002. IEEE, 2021. [6](#)
- [30] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 11966–11976. IEEE, 2022. [6](#)
- [31] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. [14](#)
- [32] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *3DV*, pages 565–571. IEEE Computer Society, 2016. [2](#)
- [33] Sergey P Morozov, AE Andreychenko, NA Pavlov, AV Vladzmyrskyy, NV Ledikhova, VA Gombolevskiy, Ivan A Blokhin, PB Gelezhe, AV Gonchar, and V Yu Chernina. Mosmeddata: Chest ct scans with covid-19 related findings dataset. *arXiv preprint arXiv:2005.06465*, 2020. [6](#), [12](#)
- [34] Mohammed Bany Muhammad and Mohammed Yeasin. Eigen-cam: Class activation map using principal components. In *IJCNN*, pages 1–7. IEEE, 2020. [2](#)
- [35] Chelsea Myers-Colet, Julien Schroeter, Douglas L. Arnold, and Tal Arbel. Heatmap regression for lesion detection using pointwise annotations. In *MILLanD@MICCAI*, pages 3–12. Springer, 2022. [2](#), [8](#)
- [36] George Papandreou, Liang-Chieh Chen, Kevin P. Murphy, and Alan L. Yuille. Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *ICCV*, pages 1742–1750. IEEE Computer Society, 2015. [2](#)
- [37] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 8024–8035, 2019. [3](#)
- [38] Pedro H. O. Pinheiro and Ronan Collobert. From image-level to pixel-level labeling with convolutional networks. In *CVPR*, pages 1713–1721. IEEE Computer Society, 2015. [2](#)
- [39] J. Radon. Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten. *Berichte über die Verhandlungen der Sächsische Akademie der Wissenschaften*, 69:262–277, 1917. [00000](#). [2](#)
- [40] Mladen Rakić, Sophie Vercauteren, Simon Van Eyndhoven, Ezequiel de la Rosa, Saurabh Jain, Sabine Van Huffel, Frederik Maes, Dirk Smeets, and Diana M. Sima. icobrain ms 5.1: Combining unsupervised and supervised approaches for improving the detection of multiple sclerosis lesions. *NeuroImage: Clinical*, 31:102707, 2021. [1](#)
- [41] Guangyu Ren, Michalis Lazarou, Jing Yuan, and Tania Stathaki. Towards automated polyp segmentation using weakly- and semi-supervised learning and deformable transformers. In *CVPR Workshops*, pages 4355–4364. IEEE, 2023. [2](#)
- [42] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. ”grabcut”: interactive foreground extraction using iterated graph cuts. *ACM Trans. Graph.*, 23(3):309–314, 2004. [2](#)
- [43] Abhijit Guha Roy, Sailesh Conjeti, Nassir Navab, Christian Wachinger, and Alzheimer’s Disease Neuroimaging Initiative. Quicknat: A fully convolutional network for quick and accurate segmentation of neuroanatomy. *NeuroImage*, 186: 713–727, 2019. [2](#)
- [44] Ashirbani Saha, Michael R. Harowicz, Lars J. Grimm, Connie E. Kim, Sujata V. Ghate, Ruth Walsh, and Maciej A. Mazurowski. A machine learning approach to radiogenomics of breast cancer: a study of 922 subjects and 529 DCE-MRI features. *British Journal of Cancer*, 119(4):508–516, 2018. [7](#), [12](#)
- [45] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017. [1](#), [2](#), [5](#)
- [46] Weixuan Sun, Jing Zhang, and Nick Barnes. 3d guided weakly supervised semantic segmentation. In *ACCV (1)*, pages 585–602. Springer, 2020. [2](#)
- [47] Shaheen Syed, Kathryn E. Anderssen, Svein Kristian Stormo, and Mathias Kranz. Weakly supervised semantic segmentation for MRI: exploring the advantages and disadvantages of class activation maps for biological image segmentation with soft boundaries. *Scientific Reports*, 13(1), 2023. [2](#)
- [48] Timothy P. Szczykutowicz, Giuseppe V. Toia, Amar Dhanantwari, and Brian Nett. A review of deep learning CT reconstruction: Concepts, limitations, and promise in clinical practice. *Current Radiology Reports*, 10(9):101–115, 2022. [2](#)
- [49] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, pages 6105–6114. PMLR, 2019. [6](#)
- [50] Zhi Tian, Chunhua Shen, Xinlong Wang, and Hao Chen. Boxinst: High-performance instance segmentation with box annotations. In *CVPR*, pages 5443–5452. Computer Vision Foundation / IEEE, 2021. [2](#)
- [51] Paul Vernaza and Manmohan Chandraker. Learning random-walk label propagation for weakly-supervised semantic segmentation. In *CVPR*, pages 2953–2961. IEEE Computer Society, 2017. [2](#)
- [52] Ostap Viniavskiy, Mariia Dobko, and Oles Dobosevych. Weakly-supervised segmentation for disease localization in chest x-ray images. In *Artificial Intelligence in Medicine*, pages 249–259. Springer International Publishing, 2020. [2](#)
- [53] Guotai Wang, Wenqi Li, Sébastien Ourselin, and Tom Vercauteren. Automatic brain tumor segmentation

- using cascaded anisotropic convolutional neural networks. In *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries - Third International Workshop, BrainLes 2017, Held in Conjunction with MICCAI 2017, Quebec City, QC, Canada, September 14, 2017, Revised Selected Papers*, pages 178–190. Springer, 2017. [2](#)
- [54] Ge Wang, Yi Zhang, Xiaojing Ye, and Xuanqin Mou. *Machine Learning for Tomographic Imaging*. IOP Publishing, 2019. [2](#)
- [55] Haofan Wang, Zifan Wang, Mengnan Du, Fan Yang, Zijian Zhang, Sirui Ding, Piotr Mardziel, and Xia Hu. Score-cam: Score-weighted visual explanations for convolutional neural networks. In *CVPR Workshops*, pages 111–119. Computer Vision Foundation / IEEE, 2020. [2](#), [7](#), [14](#)
- [56] Ross Wightman, Hugo Touvron, and Hervé Jégou. Resnet strikes back: An improved training procedure in timm. *arXiv preprint arXiv:2110.00476*, 2021. [6](#)
- [57] Ross Wightman, Haoyu Wang, Ziheng Zhang, Piotr Dollar, Yunhao Chen, and Trevor Darrell. Pytorch image models: A pytorch hub of pre-trained image classification models, 2022. [6](#)
- [58] Yuxin Wu and Kaiming He. Group normalization. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XIII*, pages 3–19. Springer, 2018. [6](#)
- [59] Yingda Xia, Lingxi Xie, Fengze Liu, Zhuotun Zhu, Elliot K. Fishman, and Alan L. Yuille. Bridging the gap between 2d and 3d organ segmentation with volumetric fusion net. In *Medical Image Computing and Computer Assisted Intervention - MICCAI 2018 - 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part IV*, pages 445–453. Springer, 2018. [2](#)
- [60] Jihye Yun, Jinkon Park, Donghoon Yu, Jaeyoun Yi, Minho Lee, Hee Jun Park, June-Goo Lee, Joon Beom Seo, and Namkug Kim. Improvement of fully automated airway segmentation on volumetric computed tomographic images using a 2.5 dimensional convolutional neural net. *Medical Image Anal.*, 51:13–20, 2019. [2](#)
- [61] Hongrun Zhang, Liam Burrows, Yanda Meng, Declan Sculthorpe, Abhik Mukherjee, Sarah E. Coupland, Ke Chen, and Yalin Zheng. Weakly supervised segmentation with point annotations for histopathology images via contrast-based variational model. In *CVPR*, pages 15630–15640. IEEE, 2023. [2](#)
- [62] Ke Zhang and Xiahai Zhuang. Cyclemix: A holistic strategy for medical image segmentation from scribble supervision. In *CVPR*, pages 11646–11655. IEEE, 2022. [2](#)
- [63] Zhou Zheng, Yuichiro Hayashi, Masahiro Oda, Takayuki Kitasaka, and Kensaku Mori. Trimix: A general framework for medical image segmentation from limited supervision. In *ACCV (6)*, pages 185–202. Springer, 2022. [2](#)
- [64] Bolei Zhou, Aditya Khosla, Àgata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 2921–2929. IEEE Computer Society, 2016. [2](#)
- [65] Xiaoyu Zhu, Jeffrey Chen, Xiangrui Zeng, Junwei Liang, Chengqi Li, Sinuo Liu, Sima Behpour, and Min Xu. Weakly supervised 3d semantic segmentation using cross-image consensus and inter-voxel affinity relations. In *ICCV*, pages 2814–2824. IEEE, 2021. [2](#)