

Unsupervised Microvascular Image Segmentation Using an Active Contours Mimicking Neural Network

Shir Gur¹, Lior Wolf^{1,2}, Lior Golgher^{3,4}, and Pablo Blinder^{3,4}

¹The School of Computer Science, Tel Aviv University

²Facebook AI Research

³School of Neurobiology, Biochemistry & Biophysics, Tel Aviv University

⁴Sagol School of Neuroscience, Tel-Aviv University

Abstract

The task of blood vessel segmentation in microscopy images is crucial for many diagnostic and research applications. However, vessels can look vastly different, depending on the transient imaging conditions, and collecting data for supervised training is laborious.

We present a novel deep learning method for unsupervised segmentation of blood vessels. The method is inspired by the field of active contours and we introduce a new loss term, which is based on the morphological Active Contours Without Edges (ACWE) optimization method. The role of the morphological operators is played by novel pooling layers that are incorporated to the network's architecture.

We demonstrate the challenges that are faced by previous supervised learning solutions, when the imaging conditions shift. Our unsupervised method is able to outperform such previous methods in both the labeled dataset, and when applied to similar but different datasets. Our code, as well as efficient pytorch reimplementations of the baseline methods VesselNN and DeepVess is available on GitHub <https://github.com/shirgur/UMIS>

1. Introduction

The field of microscopic imaging of blood vessels is evolving rapidly. For example, several emerging techniques allow rapid volumetric imaging of vascular dynamics in optically clear [5] and turbid [20, 12] living brains. Such technologies can lead to breakthroughs, both in understanding neuro-vascular interactions in the neocortex and in fine-grained medical diagnosis. However, this potential can only materialize when automatic segmentation algorithms capable of fast and accurate estimation of vascular dynamics from large volumetric movies become available.

In this work, we propose a novel unsupervised algorithm for blood vessel segmentation. The need for an unsupervised

technique arises from the limitations of the currently available training data. These limitations include: (i) the datasets size, due to the amount of expert labor required, (ii) the datasets do not represent the very high variability that exists in imaging conditions between microscopes, imaged samples, and along the same experiment. (iii) High-throughput imaging modalities[5, 12, 20] necessitate faster segmentation algorithms that do not rely on human curation.

Since blood vessels and microvessels have a well-defined tree shape, and are typically brighter than their surroundings (due to the contrast agent), classical (non-learning) algorithms can be used to tackle the problem. One relatively successful method is the Active Contours Without Edges (ACWE) method. However, the method is not accurate enough to outperform learning-based techniques.

We build an unsupervised deep learning technique that is inspired by the morphological ACWE method. The main loss that we employ is motivated by the energy function of ACWE. In addition, we introduce layers that implement the morphological operators that ACWE employs. The result is an unsupervised method that outperforms the supervised methods that were applied in this field.

As an unsupervised method, our models are less likely than the supervised methods to overfit the specific training dataset. We are able to present evidence that applied across datasets (trained on one dataset, applied to another), the new method has an even larger performance gap from the supervised learning methods in the literature.

2. Related Work

Microvascular segmentation Multi-photon laser scanning microscopy (MPLSM) allows minimally-invasive investigation of turbid samples, such as living brains, with sub-cellular resolutions [7, 35, 14, 32]. Rapid volumetric movies can be reconstructed by steering the laser beam across the volume and attributing the collected photons to the illuminated voxels [12, 20], thereby utilizing photons

that underwent multiple scatterings through the brain, from the focal volume to the detector [14].

Blood vessel segmentation has been researched in several domains, such as 3D CT and MRI [21, 8, 1], or 2D retinal blood vessels, see [28, 22, 9] to name but a few. More relevant to this paper, are the works of Tsai *et al.* [31] for neurons and microvascular segmentation and Mille *et al.* [26] for 2D and 3D branching structures extraction.

Neural network models for vessel segmentation are faster and more accurate than classical methods. Techniques for retinal blood vessel segmentation, such as [33, 9, 23] use CNNs and RNNs to perform supervised patch based segmentation of 2D images. In the images domain of two-photons microscopy, Cicek *et al.* [6] proposed a 3D-Unet for vascular segmentation, Teikari *et al.* [30] introduced VesselNN, which is a 2D-3D network architecture for 3D segmentation. Haft-Javaherian *et al.* [11] recently proposed DeepVess, which stands as the current SOTA.

Active contours First introduced by Kass *et al.* [16], active contours or snakes, are energy-minimizing methods guided by external constraint forces that pull the snake or contour towards features, such as lines and edges. These methods typically require a user-specified initial contour to start with. Variants of this method have increased the robustness and extended the applicability to new domains. A geometric model for active contours was introduced by Caselles *et al.* [2] and Yezzi *et al.* [34] with application to CT and MRI images. Kichenassamy *et al.* [17] incorporated gradients flows into the method. Geodesic Active Contour (GAC) by Caselles *et al.* [3] deform according to an intrinsic geometric measures of the image.

In the field of optimization based active contours, the ACWE work of Chan and Vese [4] and the work of Marquez-Neila *et al.* on morphological active contours [25] are the most relevant to our work. In ACWE, the borders do not have to be well-defined by gradients, and the minimization of the energy functional can be seen as a minimal partition problem. In morphological active contours, the curve evolution over time is calculated by computationally efficient morphological operators. We give a brief review of morphological ACWE in Sec. 3.

In the learning-based field, active contours did not fully integrate as a learning concept, but as a method that can be better optimized by a neural network. Rupprecht *et al.* [29] for example, trained a class-specific CNN, which predicts a vector pointing from the evolving contour towards the closest point on the boundary of the object of interest. Similarly, Marcos *et al.* [24] learn an active contour model parameterizations per instance, using a CNN. Other level-sets methods, such as [15, 18], use an additional loss that minimizes the level-set functional, while we use the Euler-Lagrange in Eq. 7 and a loss inspired by the ACWE algorithm. In all methods, supervised learning is used.



Figure 1: Illustration of the nine 3D structuring elements of \mathcal{B} . Also used as masks in the morphological pooling layer.

We introduce the first full integration between an active contour model and a learning-based model, where we use the image attachment term of the ACWE as a loss, and the morphological curvature operators as a new type of network layer. Together, we obtain an unsupervised learning algorithm for performing vascular segmentation.

3. Active Contours Without Edges

In the following section, we present a brief review of the morphological ACWE method, since it motivates our learning based method. For a detailed explanation, we kindly refer the reader to [25].

Let $\mathcal{C} : \mathbb{R}^+ \times [0, 1] \rightarrow \mathbb{R}^3$ be the parametrized 3D curve over time. We start by describing the contour evolution by a partial differential equation (PDE). A differential operator L defines the curve evolution over time with the PDE $\mathcal{C}_t = L(\mathcal{C})$, and L can be rewritten as $L(\mathcal{C}) = \mathcal{F} \cdot \mathcal{N}$, which is the product between the normal \mathcal{N} to the curve, and the scalar field \mathcal{F} , which determines the velocity of evolution for each point in the curve.

We consider a level-set representation of \mathcal{C} , defined as $u : \mathbb{R}^+ \times \mathbb{R}^3 \rightarrow \mathbb{R}$ such that $\mathcal{C}(t) = \{(x, y, z) : u(t, (x, y, z)) = 0\}$, to express the evolution of the curve as in [27, 19]:

$$\frac{\partial u}{\partial t} = \mathcal{F} \cdot |\nabla u| \quad (1)$$

where \mathcal{F} can be defined as the normal to the curve, *i.e.* $\mathcal{F} \in \{1, -1\}$, or as the intrinsic heat equation [10], *i.e.* $\mathcal{F} = \mathcal{K}$, where \mathcal{K} is the Euclidean curvature of \mathcal{C} .

3D Curvature morphological operator Monotone contrast-invariant and translation-invariant operators are called morphological operators, such as dilation and erosion. Let \mathcal{B} be a set of structuring elements that uniquely defines a morphological operator, and h is the size of that operator, we now define two morphological operators:

$$SIu(x) = \sup_{B \in \mathcal{B}} \inf_{y \in x+hB} u(y) \quad (2)$$

$$ISu(x) = \inf_{B \in \mathcal{B}} \sup_{y \in x+hB} u(y) \quad (3)$$

where \mathcal{B} contains all hyper-disks of radius 1 centered at the origin, and in our three-dimensional discrete case, the nine rectangular surfaces centered at the origin of a $3 \times 3 \times 3$ cube (see Fig. 1).

In their study, [25] showed that \mathcal{K} can be expressed by a curvature morphological operator defined as $SI \circ IS$.

Morphological ACWE Let \mathcal{C} be the parametrized curve evolution and I be an image, the Morphological ACWE functional to be minimized as defined in the work of [25, 4] is as follows:

$$\begin{aligned} F(c_1, c_2, \mathcal{C}) = & \mu \cdot \text{length}(\mathcal{C}) + v \cdot \text{area}(\text{inside}(\mathcal{C})) \\ & + \alpha \int_{\text{inside}(\mathcal{C})} \|I(x) - c_1\| dx \\ & + \beta \int_{\text{outside}(\mathcal{C})} \|I(x) - c_2\| dx \end{aligned} \quad (4)$$

where c_1 and c_2 are the mean of the values inside and outside \mathcal{C} , and defined as follows:

$$c_1(\mathcal{C}) = \frac{\int_{\text{inside}(\mathcal{C})} I(x) dx}{\int_{\text{inside}(\mathcal{C})} dx} \quad (5)$$

$$c_2(\mathcal{C}) = \frac{\int_{\text{outside}(\mathcal{C})} I(x) dx}{\int_{\text{outside}(\mathcal{C})} dx} \quad (6)$$

The parameters μ and v control the length (or curvature) and area of \mathcal{C} , as defined in [4]. The Euler-Lagrange equation for the functional F (see [25] for details) is as follows:

$$\frac{\partial u}{\partial t} = |\nabla u| \left(\mu \text{div} \left(\frac{\nabla u}{|\nabla u|} \right) - v - \alpha(I - c_1)^2 + \beta(I - c_2)^2 \right) \quad (7)$$

where in the experiments of [25], v is set to 0, and the parameters α and β are set to 1 and 2, respectively. ∇u is computed using central differences along the three axis x, y , and z . The resulting ACWE algorithm is written as follows:

$$\Gamma = |\nabla u|(\alpha(I - c_1)^2 - \beta(I - c_2)^2) \quad (8)$$

$$u^{n+\frac{1}{2}}(x) = \begin{cases} 1 & \text{if } \Gamma < 0 \\ 0 & \text{if } \Gamma > 0 \\ u^n & \text{otherwise} \end{cases} \quad (9)$$

$$u^{n+1}(x) = ((SI \circ IS)^\mu u^{n+\frac{1}{2}})(x) \quad (10)$$

where Γ is the image attachment term, as defined by [25]. The superscript is the iteration number.

4. Method

In order to implement the ACWE principle as a network, we turn the iterative energy minimization that occurs in Eq. 9 into a loss, and the morphological operations in Eq. 10 into novel morphological layers.

The segmentation network receives an input $I \in [0, 1]^{1 \times k \times m \times n}$, which is a 3D intensity-response input volume, where $k \times m \times n$ are the volume dimensions of a single intensity channel. The network outputs a segmentation map, $S \in [0, 1]^{1 \times k \times m \times n}$, and thresholding is performed to obtain the final result. To lower the ambiguity of S , and the sensitivity to thresholding, we employ a compound loss, which encourages output results to have 0 or 1 values.

4.1. Network Architecture

The network's architecture is illustrated in Fig. 2 and consists of a main Encoder-Decoder branch with skip connections, denoted as E and D_{seg} (for segmentation), respectively, followed by successive operations of Morphological Pooling Layer (Eq. 14-15) for smoothing. It is trained in an unsupervised manner, using an auxiliary reconstruction loss, provided by an additional decoder D_{rec} , which is used only during training, and outputs an estimated input \bar{I} .

The network components are then rewritten as follows:

$$\bar{S}(I) := D_{seg}(E(I)) \quad (11)$$

$$S(I) := \underbrace{SI(IS(\dots(SI(IS(\bar{S}(I))))))}_{SI \circ IS \text{ } \mu \text{ times}} \quad (12)$$

$$\bar{I} := D_{rec}(E(I)) \quad (13)$$

where \bar{S} is the segmentation before smoothing, and S is the segmentation mask obtained after applying the morphological pooling layers SI and IS in Eq. 12 μ times (the two layers are defined below).

Our Encoder architecture is based on ResNet34 [13] with 3D convolutions. In addition to the ResNet's output $E(I)$, there are two intermediate layers that are fed directly into the subsequent decoder: $C2$ ($C3$) is the output of the ResNet's 2nd (3rd) block. Each of the two decoders D_{seg} and D_{rec} consists of three Upsampling blocks with skip connections, as detailed in the appendix.

Morphological Pooling Layer Illustrated in Fig. 3 for the 2D case, the layer implements the IS and SI operations as differentiable layers, which are not learned, but play a role in the backpropagation. These layers employ the set of nine structuring elements \mathcal{B} , as explained in Sec. 3, where each element $B \in \mathcal{B}$ is a binary mask of size $3 \times 3 \times 3$ as illustrated in Fig. 1. The layers first perform masked max pooling $\forall B \in \mathcal{B}$, and then take the *maximum* or *minimum* across all results, according to the desired operation (SI or IS respectively). Formally:

$$SI(x) := \max_{B \in \mathcal{B}} -\text{MaskPool}(-x, B) \quad (14)$$

$$IS(x) := \min_{B \in \mathcal{B}} \text{MaskPool}(x, B) \quad (15)$$

$$\text{MaskPool}(x, B) := \max\{x \otimes B\} \quad (16)$$

The MaskPool function first applies an element-wise multiplication between the mask and the input, denoted by \otimes , and then takes the maximum over all locations. Fig. 3 illustrates the 2D case.

The two types of Morphological Pooling Layers are then used to smooth the intermediate segmentation \bar{S} and obtain the final segmentation mask S , according to Eq. 12, where the two morphological operations are applied $\mu = 3$ times.

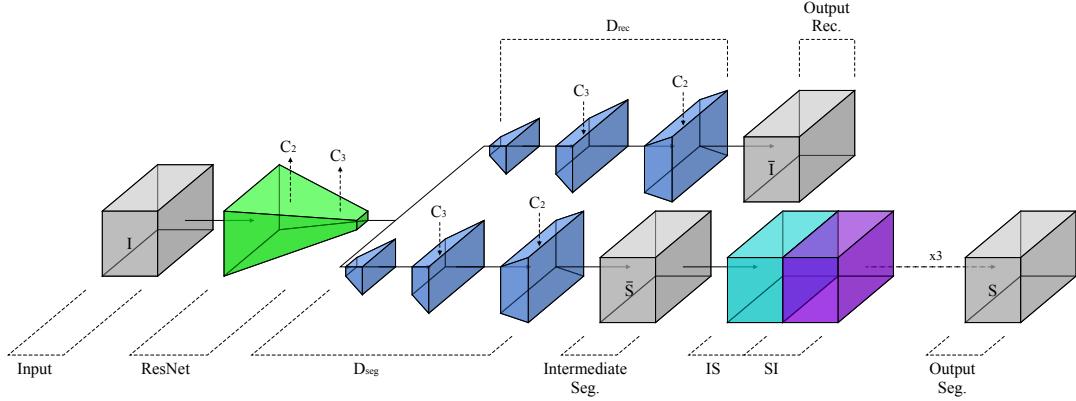


Figure 2: 3D segmentation network architecture. $C2$ and $C3$ are intermediate convolution blocks from ResNet, serving as skip-inputs for both D_{seg} and D_{rec} .

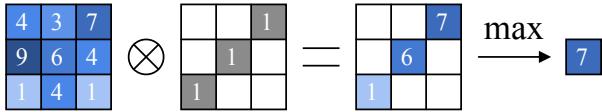


Figure 3: Illustration of 2D Masked Pooling Layer activation on a single window.

4.2. Training loss

We employ a compound loss:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{AC} + \lambda_2 \mathcal{L}_{rank} + \lambda_3 \mathcal{L}_{tight} + \lambda_4 \mathcal{L}_{rec} + \lambda_5 \mathcal{L}_{MV} + \lambda_6 \mathcal{L}_{ME} \quad (17)$$

where $\lambda_1, \dots, \lambda_6$ are weights. The various terms of the loss function are defined below.

Active contour loss Our main loss term, \mathcal{L}_{AC} , is derived from the ACWE algorithm, specifically Eq. 8 which determine the segmentation value in Eq. 9. We first rewrite Eq. 8 to fit our network terminology:

$$\Gamma = \|\nabla \bar{S}\|_1 (\alpha(I - c_1)^2 - \beta(I - c_2)^2) \quad (18)$$

where $\nabla \bar{S}$ is the intermediate segmentation derivatives in x , y , and z , computed using the central differences method. We define c_1 and c_2 similarly to Eq 5-6 as:

$$c_1 = \frac{\sum_p I(p)S(p)}{\sum_p S(p)} \quad (19) \quad c_2 = \frac{\sum_p I(p)(1 - S(p))}{\sum_p 1 - S(p)} \quad (20)$$

Following [25], we set $\alpha = 1$ and $\beta = 2$ for all experiments.

As can be seen from Eq. 9, when $\Gamma < 0$, we need to encourage the segmentation to output 1, and conversely, output 0, if $\Gamma > 0$. Therefore, we penalize elements of S for which this is not the case. This is done by the following

loss, which is computed for each point p in the volume:

$$\mathcal{L}_{AC}(p) = \begin{cases} \exp(\Gamma(p)S(p)) & \text{if } \Gamma(p) \leq 0 \\ \exp(-\Gamma(p)(1 - S(p))) & \text{if } \Gamma(p) > 0 \end{cases} \quad (21)$$

The volume loss is $\mathcal{L}_{AC} = \mathbb{E}_p \mathcal{L}_{AC}(p)$. The loss is high, if the exponent is applied to a value that is close to zero. This happens if the term $\Gamma(p)$ is negative and $S(p)$ is close to zero, or if $\Gamma(p)$ is positive and $S(p)$ is close to one.

Ranking loss \mathcal{L}_{rank} plays two roles in the learning process. Firstly, it enforces c_1 to represent the higher values in the input image, segmenting the reflective substance. Secondly, it encourages a larger gap between c_1 and c_2 .

$$\mathcal{L}_{rank} = \exp(c_2 - c_1) \quad (22)$$

Reconstruction loss Since learning is performed without supervision, we further constrain the encoder to preserve the information of the input data by adding the \mathcal{L}_{rec} reconstruction loss:

$$\mathcal{L}_{rec} = \mathbb{E}_p \left[(\bar{I}(p) - I(p))^2 + \|\nabla \bar{I}(p)\|_1 \right] \quad (23)$$

where $\nabla \bar{I}$ is a regularization (smoothing) term.

Minimal segmentation loss To avoid the case where the segmentation mask contains the entire image, which seems to be a stable solution when the training data is noisy, we encourage the network to output a tight (or minimal) segmentation:

$$\mathcal{L}_{tight} = A(S) = \sum_p S(p) \quad (24)$$

where $A(S)$ is the area of segmentation S .

Disjunctive loss We add two additional loss terms to push the segmentation output away from the intermediate values and towards 0 or 1. The Maximal Variance loss:

$$\mathcal{L}_{MV} = \exp(\mathbb{E}_p[S(p)^2] - \mathbb{E}[S(p)]^2) \quad (25)$$

and the the Maximal Entropy loss:

$$\mathcal{L}_{ME} = \mathbb{E}_p[-S(p) \cdot \log(S(p))] \quad (26)$$

4.3. Training the network

The same settings are used in all experiments and the hyperparameters are fixed to the following values: $\lambda_1 = 1$, $\lambda_2 = 10^{-2}$, $\lambda_3 = \lambda_4 = 10^{-3}$, $\lambda_5 = 10^{-3}$ and $\lambda_6 = 10^{-6}$. An Adam optimizer, with a learning rate of 0.0001, is used. Training is done on a single NVIDIA Titan-X 12G Pascal GPU, and takes about 20 minutes to train.

The 3D input volume of our network has a size of $32 \times 128 \times 128$, where 32 is for the z dimension. At train time, we crop the training samples to fit the network's input shape. At test time, a sliding window is applied, with a stride of $8 \times 16 \times 16$. The resulting outputs of the network are then averaged at every 3D location. Reflection padding is applied to the test images before slicing, to support a uniform sampling pattern across the entire volume.

4.4. Unsupervised fine tuning

An unsupervised method has the advantage that it can be applied to the test data without seeing the labels, i.e., one can naturally apply transductive learning and update the network based on the test data. This is especially helpful in the cross-domain scenario.

When we follow this protocol (we mark it explicitly as FT), the network, which was previously pre-trained on the training data, continues to train on the test data. In all of our transductive learning experiments, we use the same learning rate for both training and fine-tuning. For a fair comparison with supervised methods, we perform the additional training for the same length of time, as it takes to evaluate the test set with the relatively fast DeepVess method [11].

5. Experiments

We present three types of experiments. First, we compare our method with supervised literature neural network methods such as DeepVess [11] by Haft-Javaherian *et al.* and VesselNN [30] by Teikari *et al.*, and classic optimization methods, such as the morphological ACWE method presented above. This is done in the supervised setting, where one trains on the training split and tests on the test split of the same dataset. Second, we test the generalization capability of each method in a cross-datasets evaluation, where each method is first trained on the training set of dataset *A* and then tested on the test set of dataset *B*. Finally, we present qualitative results on a new 4D microvascular intravital dataset of neuronal activity and vascular cross-section volumetric movie. This dataset suffers from a relatively low Signal to Noise Ratio (SNR), and from sparse images at every time frame.

In our experiments we use two re-implementations of state-of-the-art supervised methods for vessels segmentation, verifying that we are able to obtain the same level of results for these baselines as the original, less efficient,

implementations. For both datasets, we compare the results of our algorithm to those of two previous unsupervised segmentation algorithms - morphological ACWE [25] and VIDA [31], in order to add unsupervised baseline methods.

5.1. Datasets

The three datasets employed consist of microvascular images produced by two-photon microscopy methods for in-vivo (turbid living brains) and in-vitro (optically cleared fixated brain slices) imaging. Due to the sensitivity of such a process in terms of noise, magnification and the subject being monitored, the data can vary in scale, resolution and SNR. This strengthens the significance of cross-domain generalization capabilities in this field, and supports our cross-dataset evaluation.

All datasets are normalized by their mean and standard deviation, followed by a range stretching, such that the minimal value is 0 and the maximal is 1.

The DeepVess dataset consists of one $200 \times 256 \times 256$ ($z \times x \times y$) in-vivo vascular image volume with expert annotations, divided into train and test. 24 additional vascular images are provided without annotation. For a fair comparison, we did not use the latter in our experiments.

The VesselNN dataset consists of 12 $20 \times 256 \times 256$ mouse cortex and tumor vascular image volume with ground truth annotations, divided into 10 train and 2 test volumes.

4D-NVIV (Neurovascular Intravital Volumetric) dataset tracks neuronal activity and vascular cross-section across an entire volume of a living mouse brain. It consists of a 4D movie spanning $110 \times 512 \times 108 \times 67527$ voxels. As done in previous work, we segment the 3D data obtained after integrating over time.

Evaluation metrics We evaluate our results using the following common segmentation metrics. Average Precision (AP) is the average precision extracted from the precision-recall graph. $F1 = 2 \cdot (precision \cdot recall) / (precision + recall)$. Sensitivity = $TP / (TP + FN)$, Specificity = $TN / (TN + FP)$, TN and FP being the true positive rate and the false positive rate, respectively. Jaccard index (JI) = $TP / (TP + FP + FN)$. DICE = $(2 \cdot TP) / (2 \cdot TP + FP + FN)$. Mean Intersection Over Union (mIoU) is the mean of all JI calculated for a sequence of thresholds between 0-1.

5.2. Results

In Tab. 1 and 2 we show results for the two annotated datasets, evaluated using common segmentation metrics. As can be seen, our method outperforms all methods in the majority of metrics, where due to the tradeoff between sensitivity and specificity, the F1 score and the Average Precision are the more informative metrics.

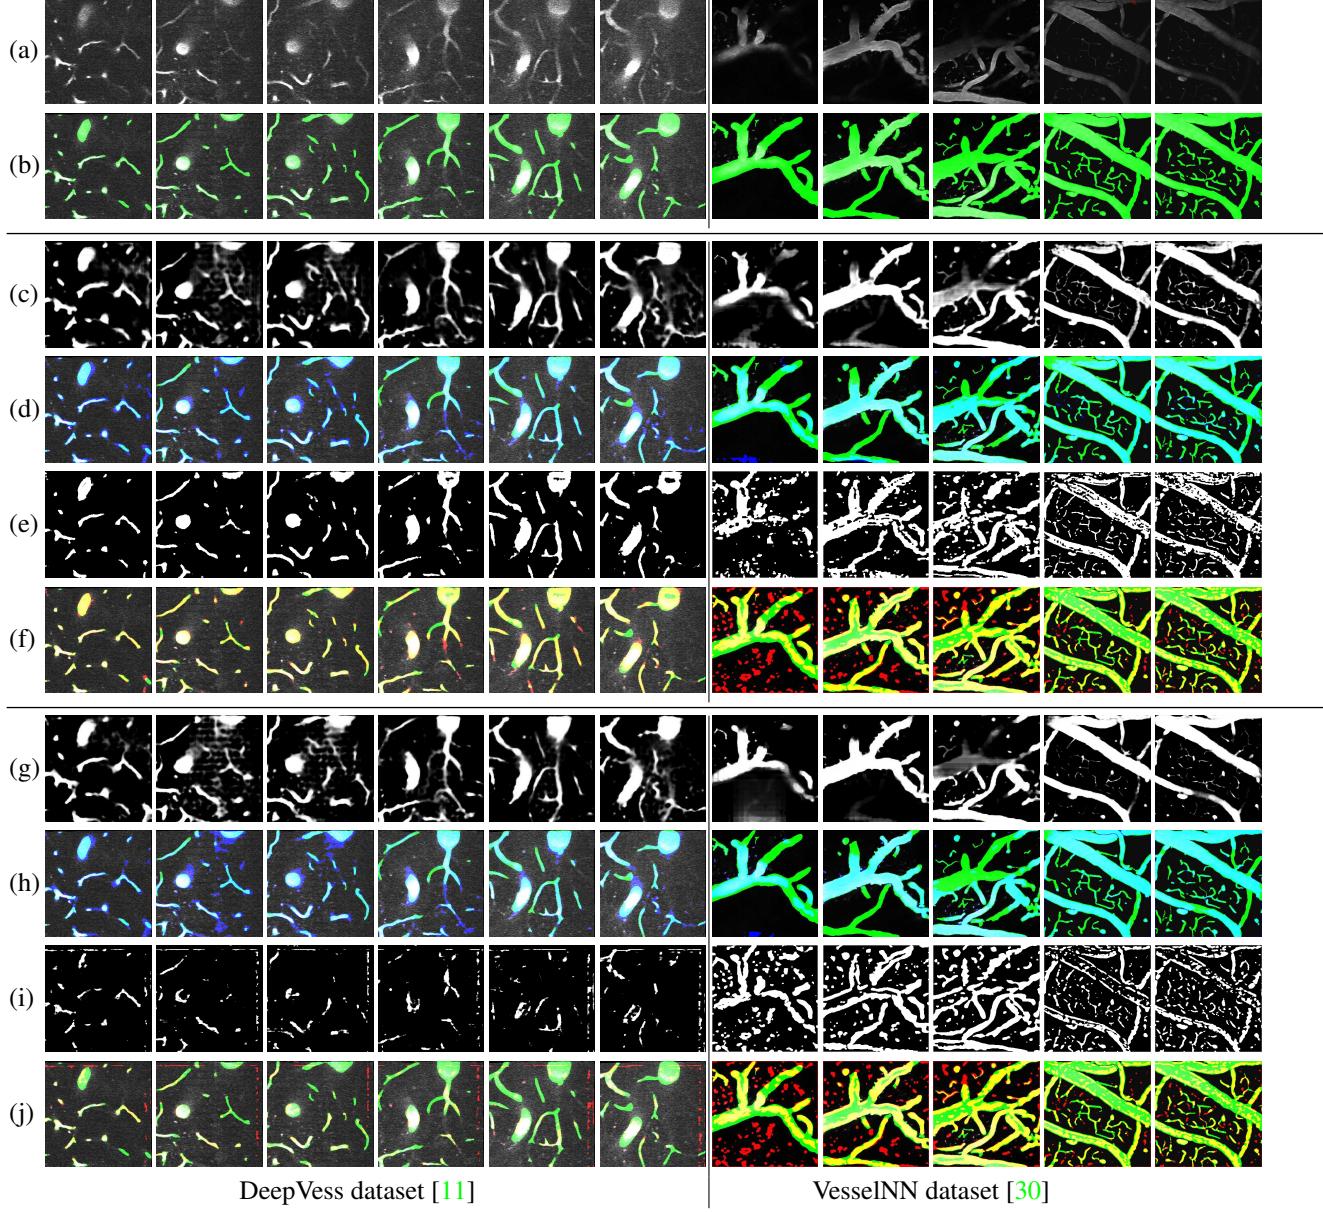


Figure 4: Qualitative segmentation results on the DeepVess dataset (left) and the VesselNN dataset (right). All images are a z slice from either dataset. **Green** - Expert annotations, **Blue** - Our segmentation, **Red** - DeepVess segmentation. **(a)** Input image **(b)** Expert annotations. **(c)-(f)** Results generated by training and testing on the same dataset, **(c)** Our output **(d)** Our + annotations, **(e)** DeepVess output, **(f)** DeepVess + annotations. **(g)-(j)** Cross-datasets results where training is performed on the opposite dataset, **(g)** Our output **(h)** Our + annotations, **(i)** DeepVess output, **(j)** DeepVess + annotations.

Cross-Datasets evaluation In this setup, each method is trained on dataset A and evaluated on dataset B . As can be seen from Tab. 3, our method significantly outperforms supervised methods. We also provide results for the fine-tuning technique, which seem to contribute most in the case of transitioning from the DeepVess to VesselNN dataset. Since VesselNN is more diverse and relatively larger, unsupervised fine-tuning is less effective in the other direction.

Qualitative results from both the within dataset and cross dataset experiments are shown in Fig. 4. Both our method and the supervised DeepVess method present a good overlap with the ground truth data on the within dataset evaluation of the DeepVess data. On the second dataset, our results are markedly better for this protocol. On the cross dataset experiment, the gap in performance in support of our method is sizable in both datasets.

Algorithm	Supervised	AP	F1	Sensitivity	Specificity	JI	DICE	mIoU
Cicek <i>et al.</i> [6]	✓	-	-	0.700	0.982	0.594	0.726	-
VesselNN [30] [†]	✓	0.867	0.732	0.868	0.989	0.777	0.732	0.765
DeepVess [11] [†]	✓	0.889	0.820	0.951	0.984	0.807	0.820	0.828
VIDA [31]		0.179	0.304	0.998	0.564	0.179	0.304	0.372
Morph-ACWE [25]		0.483	0.676	0.743	0.957	0.511	0.676	0.722
Ours		0.909	0.829	0.992	0.986	0.811	0.829	0.838

Table 1: Training and testing on the DeepVess dataset. [†]represents an improved re-implementation of the method.

Algorithm	Supervised	AP	F1	Sensitivity	Specificity	JI	DICE	mIoU
VesselNN [30] [†]	✓	0.786	0.739	0.898	0.934	0.705	0.739	0.503
DeepVess [11] [†]	✓	0.804	0.757	0.901	0.986	0.713	0.757	0.629
VIDA [31]		0.591	0.589	0.452	0.984	0.418	0.589	0.624
Morph-ACWE [25]		0.505	0.528	0.367	0.998	0.364	0.505	0.599
Ours		0.834	0.776	0.923	0.929	0.721	0.776	0.760

Table 2: Training and testing on the VesselNN dataset. [†]represents an improved re-implementation of the method.

Algorithm	Transition	AP	F1	Sensitivity	Specificity	JI	DICE	mIoU
VesselNN [30] [†]		0.528	0.188	0.110	0.991	0.108	0.188	0.469
DeepVess [11] [†]	DeepVess dataset \rightarrow	0.664	0.562	0.469	0.960	0.395	0.562	0.604
Ours	VesselNN dataset	0.684	0.662	0.833	0.915	0.501	0.662	0.681
Ours - FT		0.752	0.715	0.901	0.940	0.563	0.715	0.727
VesselNN [30] [†]		0.164	0.175	0.142	0.948	0.096	0.175	0.482
DeepVess [11] [†]	VesselNN dataset \rightarrow	0.701	0.506	0.351	0.996	0.505	0.506	0.637
Ours	DeepVess dataset	0.895	0.803	0.988	0.917	0.671	0.803	0.821
Ours - FT		0.902	0.808	0.979	0.970	0.678	0.808	0.815

Table 3: Training on A and testing on B data (A \rightarrow B). FT - Unsupervised fine-tune on A data for t seconds, where t is the time taken for DeepVess to perform evaluation. [†]represents an improved re-implementation of the method.

4D datasets segmentation The 4D-NVIV dataset differs from previous datasets in its larger field-of-view and low SNR, due to the high acquisition rate. We ran both our method, and the DeepVess method, after training on the larger and more diverse dataset of VesselNN.

The data is completely unannotated, and evaluation was done by experts who examined the output and preferred, in all cases, our output over that obtained by DeepVess. To illustrate the preference, the experts have tracked penetrating vessels across the 3D stack. This was done for four blood vessels, marked by red, green, blue and yellow squares in Fig. 5. It is evident that our method is able to produce a more reliable segmentation, where the vessel topology can be observed across cortical layers. For example, the red vessel is a penetrating artery, from which a smaller arteriole branches off at a depth of $z = 100\mu\text{m}$ below cortical surface. The ability to segment vessels in 4D data enables the highly sought after analysis of cerebral blood flow in dynamic, large field of view, imaging.

Speed Performance As can be seen from the results in Tab. 1- 3, and Fig. 5- 4, our method outperforms previous methods. We further look at the computational time that each method needs to process results. Because each method outputs a different output shape, we report the computation time in units of seconds per voxel.

The F1-score results on the DeepVess dataset is used as the other axis, and are shown for both within and cross dataset experiments. As can be seen in Fig. 7, our method outperforms the previous work along both axes.

Ablation Study We examined the significance of the different loss terms used in our method, by removing these terms one by one. Specifically, the following variants were tests: (i) without \mathcal{L}_{rank} , in which case the disjunctive loss can collapse the output of all zero or all one. (ii) without \mathcal{L}_{AC} , in which case the optimization becomes a variant of adaptive thresholding. (iii) replacing S with \bar{S} in Eq. 21, examining the effect of the morphological pooling layer. (iv)

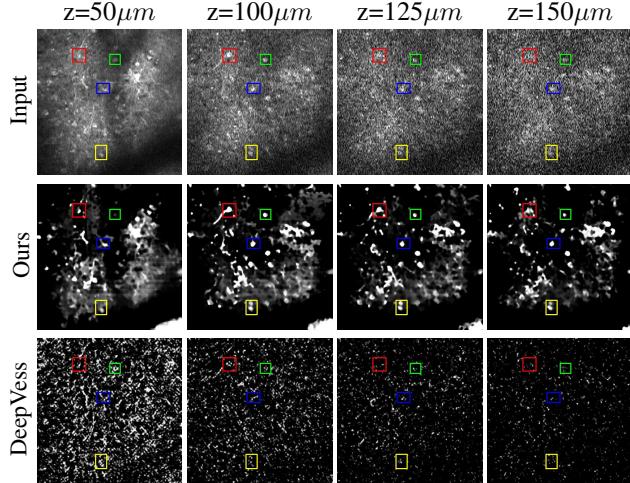


Figure 5: Qualitative results from a 4D movie. **Columns** - depth from brain surface. The SNR decreases with imaging depth. **Rows** - Input, Our output, DeepVess output. **Colored boxes** - expert markers of penetrating vessels.

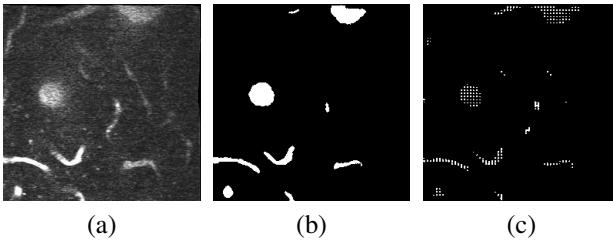


Figure 6: Illustration of the effect caused by replacing S with \bar{S} in Eq. 21. (a) the input image. (b) the output of our method. (c) the output of our method when not employing the morphological layers.

using only \mathcal{L}_{AC} , which is the new loss term we introduce. (v) using only $\mathcal{L}_{AC} + \mathcal{L}_{rank}$, which are the two main losses.

As can be seen in Tab. 4, the role of both \mathcal{L}_{AC} and \mathcal{L}_{rank} is very significant. One can note that the use of \mathcal{L}_{AC} (Eq. 21) as a single loss already produces favorable results, supporting this novel loss term. The importance of the morphological layers can be observed by the drop in performance, when replacing S by \bar{S} in Eq. 21. Visually, as can be seen in Fig. 6, this results in a stripped output. We hypothesize that this pattern emerges as a way to increase the local gradient in Eq. 18.

6. Conclusion

The quantitative analysis of blood vessel microscopy images requires accurate segmentation capabilities. While the task can be solved in a supervised way, the lack of carefully curated training data and the large domain shift between ex-

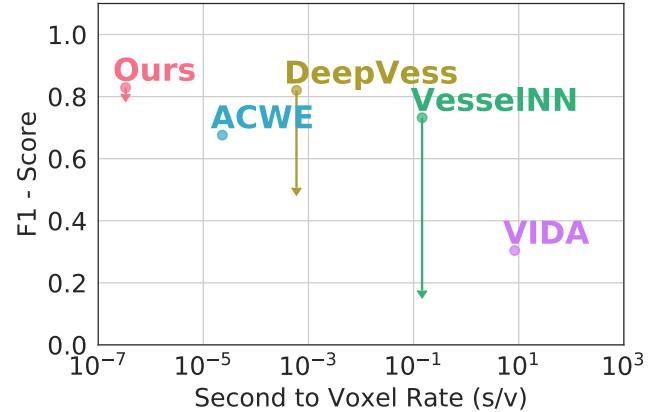


Figure 7: Computation rate vs. F1-score on DeepVess dataset. Vector lines represent the drop in performance, when performing cross-dataset evaluation.

Losses	AP	F1	JI	DICE	mIoU
\mathcal{L} w/o \mathcal{L}_{rank}	0.100	0.159	0.086	0.159	0.459
\mathcal{L} w/o \mathcal{L}_{AC}	0.617	0.437	0.280	0.437	0.608
Eq. 21 w/ \bar{S}	0.831	0.707	0.618	0.707	0.779
\mathcal{L}_{AC}	0.866	0.798	0.665	0.798	0.805
$\mathcal{L}_{AC} + \mathcal{L}_{rank}$	0.889	0.817	0.689	0.817	0.821
Ours \mathcal{L}	0.909	0.829	0.708	0.829	0.838

Table 4: Ablation study showing results for various losses.

perimental settings, create the need for alternative methods.

The morphological ACWE method is a relatively recent active contour algorithm that is well-motivated and robust. In our work, it is reincarnated as a deep learning method that is able to benefit from employing a training set, as well as from the accumulated experience in our community of effectively employing CNNs. Our work, therefore, provides not just a state of the art solution that is based on novel losses and new types of layers, but also a case study for turning classical and powerful computer vision techniques into deep learning methods.

Acknowledgement

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (ERC grants CoG 725974 and StG 639416). The contribution of the first author is part of a Ph.D. thesis research conducted at Tel Aviv University. The authors thank David Kain for conducting the mouse surgery and Yulia Mitiagin for segmenting the said datasets using the VIDA suite [31].

References

[1] Pierre Bouillot, Bénédicte M. A. Delattre, Olivier Brina, Rafik Ouared, Mohamed Farhat, C Chnafa, David A. Steinman, Karl olf Lovblad, Vitor Mendes Pereira, and Maria Isabel Vargas. 3d phase contrast mri: Partial volume correction for robust blood flow quantification in small intracranial vessels. *Magnetic resonance in medicine*, 79(1):129–140, 2018. [2](#)

[2] Vicent Caselles, Francine Catté, Tomeu Coll, and Françoise Dibos. A geometric model for active contours in image processing. *Numerische mathematik*, 66(1):1–31, 1993. [2](#)

[3] Vicent Caselles, Ron Kimmel, and Guillermo Sapiro. Geodesic active contours. *International journal of computer vision*, 22(1):61–79, 1997. [2](#)

[4] Tony F Chan and Luminita A Vese. Active contours without edges. *IEEE Transactions on image processing*, 10(2):266–277, 2001. [2, 3](#)

[5] Karishma Chhabria, Karen Plant, Oliver Bandmann, Robert N Wilkinson, Chris Martin, Elisabeth Kugler, Paul A Armitage, Paola LM Santoscoy, Vincent T Cunliffe, Jan Huisken, et al. The effect of hyperglycemia on neurovascular coupling and cerebrovascular patterning in zebrafish. *Journal of Cerebral Blood Flow & Metabolism*, page 0271678X18810615, 2018. [1](#)

[6] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *International conference on medical image computing and computer-assisted intervention*, pages 424–432. Springer, 2016. [2, 7](#)

[7] Winfried Denk, James H. Strickler, and Watt W. Webb. Two-photon laser scanning fluorescence microscopy. *Science*, 248(4951):73–6, 1990. [1](#)

[8] Bruce Fischl, David H. Salat, Evelina Busa, Marilyn Albert, Megan E Dieterich, Christian Haselgrove, André J. W. van der Kouwe, Ron J. Killiany, David Kennedy, Shuna Klaveness, Albert Montillo, Nikos Makris, Bruce Rosen, and Anders M. Dale. Whole brain segmentation automated labeling of neuroanatomical structures in the human brain. *Neuron*, 33:341–355, 2002. [2](#)

[9] Huazhu Fu, Yanwu Xu, Damon Wing Kee Wong, and Jiang Liu. Retinal vessel segmentation via deep learning network and fully-connected conditional random fields. *2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI)*, pages 698–701, 2016. [2](#)

[10] Frédéric Guichard, Jean-Michel Morel, and Robert Ryan. Contrast invariant image analysis and pdes. *preprint*, 35, 2004. [2](#)

[11] Mohammad Haft-Javaherian, Linjing Fang, Victorine Muse, Chris B Schaffer, Nozomi Nishimura, and Mert R Sabuncu. Deep convolutional neural networks for segmenting 3d in vivo multiphoton images of vasculature in alzheimer disease mouse models. *arXiv preprint arXiv:1801.00880*, 2018. [2, 5, 6, 7](#)

[12] Hagai Har-Gil, Lior Golgher, Shai Israel, David Kain, Ori Cheshnovsky, Moshe Parnas, and Pablo Blinder. Pysight: plug and play photon counting for fast continuous volumetric intravital microscopy. *Optica*, 5(9):1104–1112, 2018. [1](#)

[13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. [3](#)

[14] Erich E. Hoover and Jeff A. Squier. Advances in multiphoton microscopy technology. *Nature photonics*, 7(2):93–101, 2013. [1, 2](#)

[15] Ping Hu, Bing Shuai, Jun Liu, and Gang Wang. Deep level sets for salient object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2300–2309, 2017. [2](#)

[16] Michael Kass, Andrew Witkin, and Demetri Terzopoulos. Snakes: Active contour models. *International journal of computer vision*, 1(4):321–331, 1988. [2](#)

[17] Satyanad Kichenassamy, Arun Kumar, Peter Olver, Allen Tannenbaum, and Anthony Yezzi. Gradient flows and geometric active contour models. In *Proceedings of IEEE International Conference on Computer Vision*, pages 810–815. IEEE, 1995. [2](#)

[18] Youngeun Kim, Seunghyeon Kim, Taekyung Kim, and Changick Kim. Cnn-based semantic segmentation using level set loss. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1752–1760. IEEE, 2019. [2](#)

[19] Ron Kimmel. *Numerical geometry of images: Theory, algorithms, and applications*. Springer Science & Business Media, 2003. [2](#)

[20] Lingjie Kong, Jianyong Tang, Justin P Little, Yang Yu, Tim Lämmermann, Charles P Lin, Ronald N Germain, and Meng Cui. Continuous volumetric imaging via an optical phase-locked ultrasound lens. *Nature methods*, 12(8):759, 2015. [1](#)

[21] David Lesage, Elsa D. Angelini, Isabelle Bloch, and Gareth Funka-Lea. A review of 3d vessel lumen segmentation techniques: Models, features and extraction schemes. *Medical image analysis*, 13(6):819–45, 2009. [2](#)

[22] Pawex0142 Liskowski and Krzysztof Krawiec. Segmenting retinal blood vessels with deep neural networks. *IEEE Transactions on Medical Imaging*, 35:2369–2380, 2016. [2](#)

[23] Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Pablo Andrés Arbeláez, and Luc Van Gool. Deep retinal image understanding. In *MICCAI*, 2016. [2](#)

[24] Diego Marcos, Devis Tuia, Benjamin Kellenberger, Lisa Zhang, Min Bai, Renjie Liao, and Raquel Urtasun. Learning deep structured active contours end-to-end. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8877–8885, 2018. [2](#)

[25] Pablo Marquez-Neila, Luis Baumela, and Luis Alvarez. A morphological approach to curvature-based evolution of curves and surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(1):2–17, 2014. [2, 3, 4, 5, 7](#)

[26] Julien Mille and Laurent D. Cohen. Deformable tree models for 2d and 3d branching structures extraction. *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 149–156, 2009. [2](#)

- [27] Stanley Osher and James A Sethian. Fronts propagating with curvature-dependent speed: algorithms based on hamilton-jacobi formulations. *Journal of computational physics*, 79(1):12–49, 1988. [2](#)
- [28] Elisa Ricci and Renzo Perfetti. Retinal blood vessel segmentation using line operators and support vector classification. *IEEE Transactions on Medical Imaging*, 26:1357–1365, 2007. [2](#)
- [29] Christian Rupprecht, Elizabeth Huaroc, Maximilian Baust, and Nassir Navab. Deep active contours. *arXiv preprint arXiv:1607.05074*, 2016. [2](#)
- [30] Petteri Teikari, Marc Santos, Charissa Poon, and Kullervo Hyynnen. Deep learning convolutional networks for multi-photon microscopy vasculature segmentation. *arXiv preprint arXiv:1606.02382*, 2016. [2, 5, 6, 7](#)
- [31] Philbert S. Tsai, John P. Kaufhold, Pablo Blinder, Beth Friedman, Patrick J. Drew, Harvey J. Karten, Patrick D. Lyden, and David Kleinfeld. Correlations of neuronal and microvascular densities in murine cortex revealed by direct counting and colocalization of nuclei and vessels. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 29 46:14553–70, 2009. [2, 5, 7, 8](#)
- [32] Alan Urban, Lior Golgher, Clément Brunner, Amos Gdalyahu, Hagai Har-Gil, David Kain, Gabriel Montaldo, Laura Sironi, and Pablo Blinder. Understanding the neurovascular unit at multiple scales: advantages and limitations of multi-photon and functional ultrasound imaging. *Advanced drug delivery reviews*, 119:73–100, 2017. [1](#)
- [33] Aaron Wu, Ziyue Xu, Mingchen Gao, Mario Buty, and Daniel J. Mollura. Deep vessel tracking: A generalized probabilistic approach via deep learning. *2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI)*, pages 1363–1367, 2016. [2](#)
- [34] Anthony Joseph Yezzi, Satyanad Kichenassamy, Arun Kumar, Peter J. Olver, and Allen R. Tannenbaum. A geometric snake model for segmentation of medical imagery. *IEEE Transactions on Medical Imaging*, 16:199–209, 1997. [2](#)
- [35] Warren R. Zipfel, Rebecca M. E. Williams, and Watt W. Webb. Nonlinear magic: multiphoton microscopy in the biosciences. *Nature Biotechnology*, 21:1369–1377, 2003. [1](#)