

Microvascular Dynamics from 4D Microscopy Using Temporal Segmentation

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

¹*School of Computer Science, Tel Aviv University*

²*Facebook AI Research*

³*School of Neurobiology, Biochemistry & Biophysics, Faculty of Life Sciences, Tel Aviv University*

⁴*Sagol School of Neuroscience, Tel-Aviv University*

Recently developed methods for rapid continuous volumetric two-photon microscopy facilitate the observation of neuronal activity in hundreds of individual neurons and changes in blood flow in adjacent blood vessels across a large volume of living brain at unprecedented spatio-temporal resolution. However, the high imaging rate necessitates fully automated image analysis, whereas tissue turbidity and photo-toxicity limitations lead to extremely sparse and noisy imagery. In this work, we extend a recently proposed deep learning volumetric blood vessel segmentation network, such that it supports temporal analysis. With this technology, we are able to track changes in cerebral blood volume over time and identify spontaneous arterial dilations that propagate towards the pial surface. This new capability is a promising step towards characterizing the hemodynamic response function upon which functional magnetic resonance imaging (fMRI) is based.

Keywords: Two-photon microscopy; Microvasculature; Segmentation; Deep learning.

1. Introduction

The mammalian neocortex is innervated by a dense, regulated network of blood vessels known as the cortical angiome.¹ The cortical angiome exhibits neurovascular coupling, namely a temporary change in cerebral blood flow triggered by neuronal activity through direct and indirect signalling pathways, replenishing the surrounding tissue with oxygen and nutrients and removing excess heat and waste.²⁻⁵ Impaired neurovascular coupling is associated with a variety of debilitating pathological conditions, such as dementia, hypertension, diabetes and Alzheimer's disease.^{3,4,6} While the structural properties of the cortical angiome have been considerably elucidated,¹ our understanding of its functional properties is limited, owing in part to the insufficient spatiotemporal resolution of existing imaging techniques.⁵ Simply put, it is still unknown how individual microvessels react to individual neuronal action potentials, with preliminary evidence suggesting that the vascular response is mostly driven by specific subtypes of interneurons.^{3,7,8}

Previous attempts to measure the vascular response to neuronal activity were limited to imaging changes in vascular diameter one plane at a time.⁷⁻¹¹ Most works have repeatedly exposed the animal to an artificial prolonged sensory stimulation over hundreds of consecutive trials, eliciting a vigorous neuronal activation that gave rise to a measurable vascular re-

sponse.^{7,9-11} By repeating these visual,⁹ olfactory¹⁰ or somatosensory^{7,11} sensory stimuli while focusing on differing layers of the cortex, a laminar difference in the average onset time and time-to-peak of the vascular response was revealed.¹² In particular, instances of vasodilation begun earlier in the deepest cortical layers, suggesting that vasodilation propagates upwards along penetrating arteries.^{7,11,12}

With neuronal activity *per se* being at the primary focus of most neuroimaging labs,¹³ several imaging methods have been recently tailored for rapidly tracking neuronal activity across considerable large brain volumes, with cellular or near-cellular resolution. These include light field microscopy,¹⁴ lensless imaging,¹⁵ scanned line angular projection microscopy,¹⁶ and reconstruction of 3D imagery from 2D images using deep neuronal networks.¹⁷ While the spatial resolution of these methods is sufficient to discern neuronal cell bodies with little cross-talk, it is insufficient for tracking minuscule changes in cerebral blood diameter, which cause considerable changes in cerebral blood flow. Conversely, while optical coherence tomography and functional ultrasound imaging allow noninvasive, label-free tracking of vascular dynamics over large fields of views, they are incapable of tracking neuronal activity with single cell resolution.⁵ The invention of the ultrasonic varifocal lens allowed axial scanning at > 100 kHz rates, enabling rapid continuous volumetric multi-photon imaging.^{18,19} It is now technically possible to track the activity of hundreds of neurons and vasoactivity along neighbouring blood vessels, simultaneously across a large brain volume.¹⁹ To name but a few of the benefits of continuous volumetric imaging over traditional planar imaging:

- (1) The propagation of vasodilation and vasoconstriction through cortical vessels can be directly observed during spontaneous brain activity, rather than indirectly deduced from averaging repeated evoked trials at differing cortical depths.
- (2) Most neuronal cell bodies surrounding a given blood vessel segment are observable with volumetric imaging, but not in planar imaging. Therefore a greater proportion of the neuronal activity that drives vasoactivity is accounted for.
- (3) Instances in which neuronal action potentials at a given cortical layer affect metabolic demand at another cortical layer can be accounted for.
- (4) Axial motion (z-drift), that is known to introduce a considerable bias during planar imaging of neuronal activity,²⁰ can be accounted for in rapid continuous volumetric imaging.¹⁸
- (5) Cerebral blood volume can be directly measured for each vessel segment, rather than derived from its diameter along an arbitrary axis in an error-prone fashion.²¹

However, the size of 4D datasets generated by volumetric imaging precludes manual time-lapse segmentation of the imaged blood vessels. Furthermore, the sparse imagery obtained by rapid volumetric two-photon imaging^{19,22} complicates time-lapse vascular segmentation even for trained annotators. The development of accurate algorithms for automated vascular segmentation in 3D-movies is therefore essential for the analysis of neurovascular interactions.

In this paper we show the potential applications of automated angiomal segmentation for the tracking of changes in cerebral blood volume over time, as well as for the identification of spontaneous vascular dilations propagating along penetrating arteries towards the pial surface.

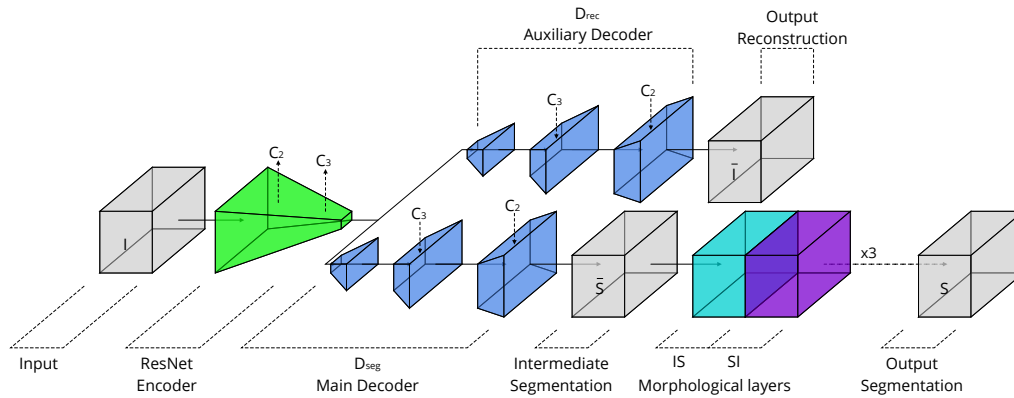


Fig. 1. Network architecture. C_2 and C_3 are intermediate activations from the ResNet encoder, which are skip-connected to both D_{seg} and D_{rec} at the marked locations.

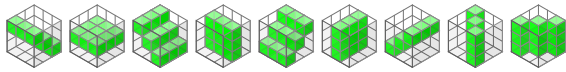


Fig. 2. Illustration of the 3D structuring elements of \mathcal{B} . The elements are used as masks in the morphological pooling layer.

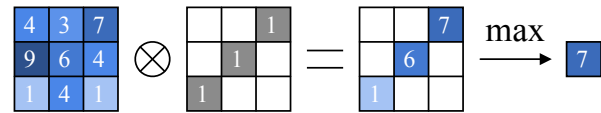


Fig. 3. A 2D Masked Pooling Layer.

To the best of our knowledge, this is the first time that such capabilities are shown.

From a technical perspective, our work extends a recently proposed deep learning technique for microvessel segmentation²³ called the ACWE network, which is presented in Section 2. This network is trained in an unsupervised manner and does not require carefully labeled training samples, in contrast to other recent deep learning approaches.^{24–28} It is based on the optimization problem minimized by the morphological Active Contours Without Edges method,²⁹ which is converted into a deep learning solution. The ACWE network was shown to outperform both classical active contour methods as well as the recent deep learning solutions, and to be robust to domain shift across datasets.²³

While the ACWE network is able to perform well on the task of extracting a single (time-collapsed) microvascular map from a given 4D volume (the same task that is being handled by other recent contributions^{23–25}), there is no prior work capable of handling 4D datasets. We demonstrate that the state of the art static map obtained by ACWE is insufficient for performing an analysis of the temporal dynamics. We therefore propose a method to obtain a sequence of such maps, that makes use of the temporal dynamics, and which is much more suitable for our analysis. The full temporal treatment is made possible by the novel skeleton layer, which ties the segmentation results of the individual frames together.

2. The ACWE network

The ACWE network²³ reincarnates the ACWE method²⁹ as a deep learning technique. This is done by replacing the iterative energy minimization that occurs in the classical method into a loss, and the morphological operations of this method into morphological layers. The 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 and 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. The network's architecture is illustrated in Fig. 1 and consists of a main Encoder-Decoder branch with skip connections, denoted as E and D_{seg} (for *segmentation*), followed by successive operations of Morphological Pooling Layer (Eq. 4-5) 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 a reconstruction \bar{I} . The network components are then rewritten as follows:

$$\bar{I} := D_{rec}(E(I)) \quad (1) \quad S(I) := \underbrace{SI(IS(\dots(SI(IS(\bar{S}(I))))))}_{SI \circ IS \ \mu \ \text{times}} \quad (3)$$

$$\bar{S}(I) := D_{seg}(E(I)) \quad (2)$$

where the operator $SI \circ IS$ repeats μ times, \bar{S} is the segmentation before smoothing, and S is the segmentation mask obtained after applying the morphological pooling layers SI and IS $\mu = 3$ times (the two layers are defined below). The Encoder architecture is based on ResNet34,³⁰ where 2D convolutions are replaced with 3D ones. Each of the two decoders D_{seg} and D_{rec} consists of three upsampling blocks with skip connections.

The morphological layers IS and SI employ a set of nine structuring elements \mathcal{B} , following,²⁹ where each element $B \in \mathcal{B}$ is a binary mask of size $3 \times 3 \times 3$ as illustrated in Fig. 2. The layers 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, the function $MaskPool(x, B) = \max\{x \otimes B\}$ first applies an element-wise multiplication between the mask and the input, denoted by \otimes , and then takes the maximum over all locations, see Fig. 3 for an illustration of the 2D case. The layers are define as:

$$SI(x) = \max_{B \in \mathcal{B}} -MaskPool(-x, B) \quad (4) \quad IS(x) = \min_{B \in \mathcal{B}} MaskPool(x, B) \quad (5)$$

The active contour loss term, \mathcal{L}_{AC} , is derived from the ACWE algorithm. Let Γ be the energy that the ACWE minimizes, defined as:

$$\Gamma = \|\nabla \bar{S}\|_1 (\alpha(I - c_1)^2 - \beta(I - c_2)^2) \quad (6)$$

where I is the input volume, and c_1 (c_2) are the average intensities inside and outside the segmentation mask S , i.e, $c_1 = \frac{\sum_p I(p)S(p)}{\sum_p S(p)}$, and $c_2 = \frac{\sum_p I(p)(1-S(p))}{\sum_p 1-S(p)}$, where p is a voxel. The loss is averaged over all 3D points, where per point in the 3D volume p it is given as:

$$\mathcal{L}_{AC}(p) = \begin{cases} \exp(\Gamma(p)S(p)) & \text{if } \Gamma(p) \leq 0 \\ \exp(-\Gamma(p)(1-S(p))) & \text{if } \Gamma(p) > 0 \end{cases} \quad (7)$$

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. Therefore, the loss minimizes Γ . The other loss terms are briefly given as:

$$\mathcal{L}_{rank} = \exp(c_2 - c_1) \quad (8)$$

$$\mathcal{L}_{rec} = \mathbb{E}_p [(\bar{I}(p) - I(p))^2 + \|\nabla \bar{I}(p)\|_1] \quad (9)$$

$$\mathcal{L}_{tight} = \sum_p S(p) \quad (10)$$

$$\mathcal{L}_{MV} = \exp(\mathbb{E}_p[S(p)^2] - \mathbb{E}[S(p)]^2) \quad (11)$$

$$\mathcal{L}_{ME} = \mathbb{E}_p[-S(p) \cdot \log(S(p))] \quad (12)$$

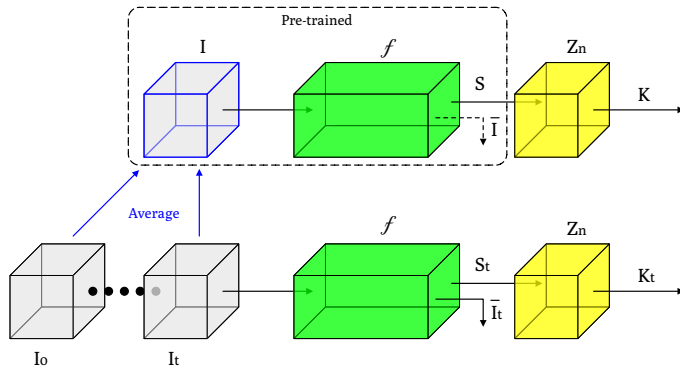


Fig. 4. Temporal network architecture. f represent the same architecture as in Fig. 1, I_t , S_t , \bar{I}_t and K_t are the temporal input, segmentation, reconstruction and skeleton. Z_n – skeletonization layer.

Eq. 8 pushes c_1 to be significantly larger than c_2 , Eq. 9 is the loss of the reconstruction pathway, where $\nabla \bar{I}$ is a smoothing term, Eq. 10 pushes the segmentation S to be minimal, and Eq. 11- 12 encourage S to have semi-binary values close to 0 or close to 1.

The ACWE network employs a compound loss ($\lambda_1, \dots, \lambda_6$ are weights):

$$\mathcal{L} = \lambda_1 \mathcal{L}_{AC} + \lambda_2 \mathcal{L}_{rank} + \lambda_3 \mathcal{L}_{tight} + \lambda_4 \mathcal{L}_{rec} + \lambda_5 \mathcal{L}_{MV} + \lambda_6 \mathcal{L}_{ME} \quad (13)$$

3. Method

Our method extends the time-collapsed segmentation,²³ where we add an additional novel skeletonization layer on top of the network and perform per-frame segmentation of the 4D movie. The resulting skeleton serves as an anchoring structure that ties all temporal results together regardless of the transient vascular changes. We first describe our novel differentiable 3D skeletonization layer and proceed with the algorithm specifics.

Skeleton layer The skeleton layer promotes spatially coherent tree-like structures. The novel iterative layer is fully differentiable with respect to the input image, and it is based on the *MaskPool* layers and the extension of Lanturjoul’s formula³¹ by Beucher *et al.*³².

Let S be the segmentation output of our network, the skeleton layer’s output is given by Z_n , and it is obtained after n iterations, starting from $Z_0 = S$:

$$erosion := \min_{B \in \mathcal{B}} -MaskPool(-x, B) \quad (14)$$

$$Z_n = (Z_{n-1} \ominus \mathcal{B}) \cup R(Z_{n-1}) \quad (16)$$

$$dilation := \max_{B \in \mathcal{B}} MaskPool(x, B) \quad (15)$$

where $R(Q) = Q/(Q)_{\mathcal{B}}$ for some input Q , and the operator $(\cdot)_{\mathcal{B}}$ denotes the *open* operator, i.e., erosion followed by dilation. The union marks a point-wise maximum, and the erosion operator is denoted by \ominus . In Fig. 5(e) we show a sample of the temporal skeleton output.

Training First, network $f = D_{seg} \circ E$ trains on the time-collapsed data obtained by averaging all time frames and generates a segmentation S . The skeleton layer is then used ($n = 5$) to produce the anchor skeleton K from S .

The temporal segmentation network is then the same f , retrained on each sparse frame of the 4D image (a 3D volume denoted I_t), with an additional skeleton loss, which encourage the temporal segmentation to be aligned with the static time-collapsed skeleton K . For each temporal segmentation S_t made by f , we compute the skeleton $K_t = Z_n|_{Z_0=S_t}$ and the loss:

$$\mathcal{L}_t = \mathcal{L} + \mathbb{E}_p \|K(p) - K_t(p)\|_1 \quad (17)$$

where the loss \mathcal{L} (Eq. 13) is computed on I_t . Fig. 4 illustrate the entire training process.

Pre/Post processing We initially downsample the 4D movie to 60Hz as the input to our network. The resulting 4D segmented movie is then smoothed by a moving average along the z axis with a width of one tenth of the volume.

4. Results

All imaging experiments and surgical procedures were approved by the Tel Aviv University ethics committee for animal use and welfare and followed pertinent Institutional Animal Care and Use Committee (IACUC) and local guidelines. Please see²² for a full description of surgical procedures. Neuronal activity was monitored in an adult male mouse from the C57BL/6J-Tg(Thy1-GCaMP6f)GP5.5Dkim/J transgenic line. Vascular dynamics were monitored by injecting Fluorescein isothiocyanate (FITC) conjugated with 2MDa dextran. Prior to imaging, the mouse was habituated to the imaging conditions across five consecutive days. To minimize motion artifacts during imaging, its head was restrained to a custom-made holder and its platform was clamped to the imaging stage.

We examine two datasets acquired using rapid volumetric two-photon laser scanning microscopy.²² The first dataset (Fig. 5(a)) tracks cerebral blood volume in nine penetrating arteries and neuronal activity in 103 adjacent neurons (see Fig. 5(a-6)), within a volume of living mouse brain spanning $430 \times 440 \times 200 \mu m^3$, imaged over 536 sec. at a rate of 125.87 volumes per second. The 4D movie was parsed into $110 \times 512 \times 108$ voxels and its first 60 sec. (7552 volumes) were selected for analysis. The second dataset (Fig. 5(b)) was acquired within the same imaging session, in the same mouse, at the same magnification, and was centered on the same field of view, albeit spanning only $92 \times 440 \times 200 \mu m^3$, imaged over 268 seconds at a rate of 125.87 volumes per second. The 4D movie was parsed into $110 \times 512 \times 108$ voxels and its first 60 seconds (7552 volumes) were selected for analysis. In both datasets, a binning of 2 over time was used for ease of computation, yielding 3776 binned volumes.

All fluorescence values were normalized by their mean and standard deviation, followed by a range stretching, such that the minimal value is 0 and the maximal is 1. A human expert has identified and annotated twelve penetrating arteries in the 3D volume, and we rely on this annotation in our analysis. The human annotation takes the form of rectangular region in each z-slice. We note that our segmentation results could promote the development of automatic vessel annotation methods. However, in this work, we focus on the much more critical bottleneck of obtaining reliable vessel activity measurements.

In Fig. 5 we present the raw data as well as the output of multiple steps of processing. These include: (1) Time-collapsed original data (2) Raw video, (3) Time-collapsed segmentation, (4) Time-varying segmentation (5) Time-varying skeleton. We have also annotated the twelve penetrating vessels. As can be seen, the raw data is almost non informative, while in our time-varying method, the vessels are most clearly visible.

The evaluation is limited by the inability to measure the dynamic behavior by other means. We therefore ran multiple auxiliary experiments as sanity check. In one experiment, we have acquired the same vessels with twice the magnification, in addition to the two 4D movies. To assess our method, we compared the diameter of each annotated vessel, in the

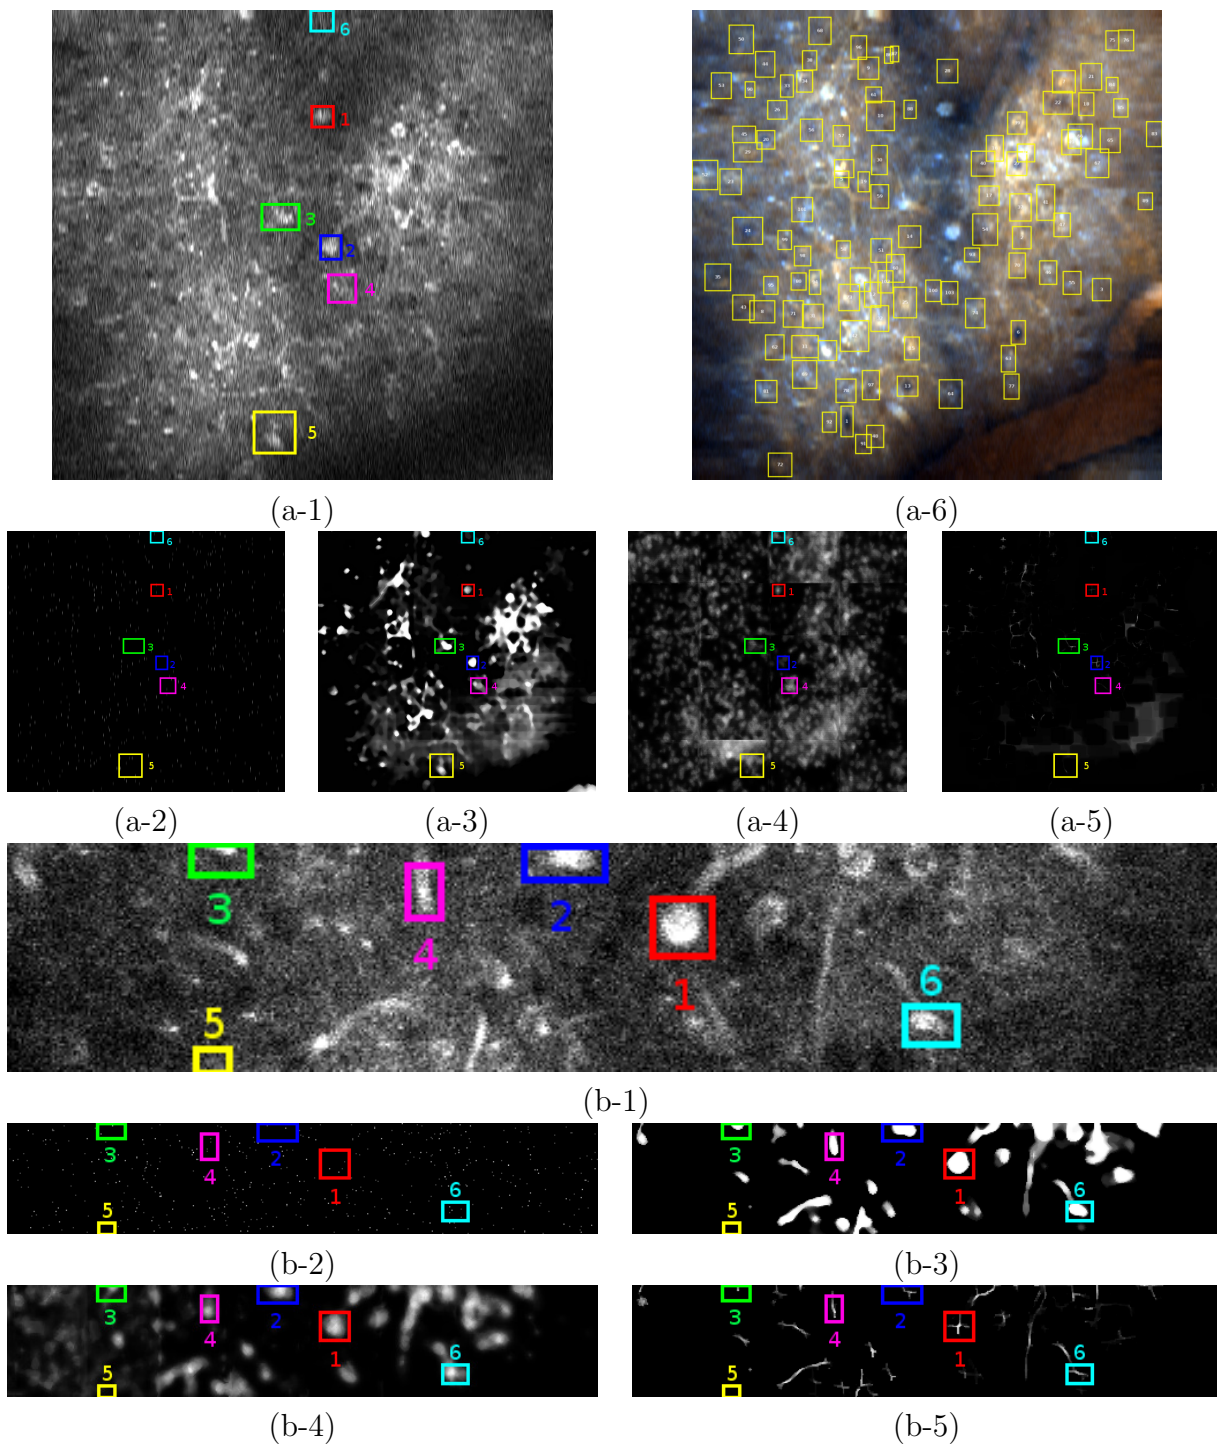


Fig. 5. Two 4D movies, (a-*) and (b-*), with sampled images shown for a depth of $z = 75\mu m$ below pial surface, at $t=1.5$ seconds. (*-1) Time-collapsed, (*-2) Raw video Time-collapsed original data, (*-3) Time-collapsed segmentation, (*-4) Time-varying segmentation (*-5) Time-varying skeleton. Annotated vessels are marked from 1-6. (a-6) shows 103 neuronal cell bodies demarcated in a depth-color-coded projection of the same volume.

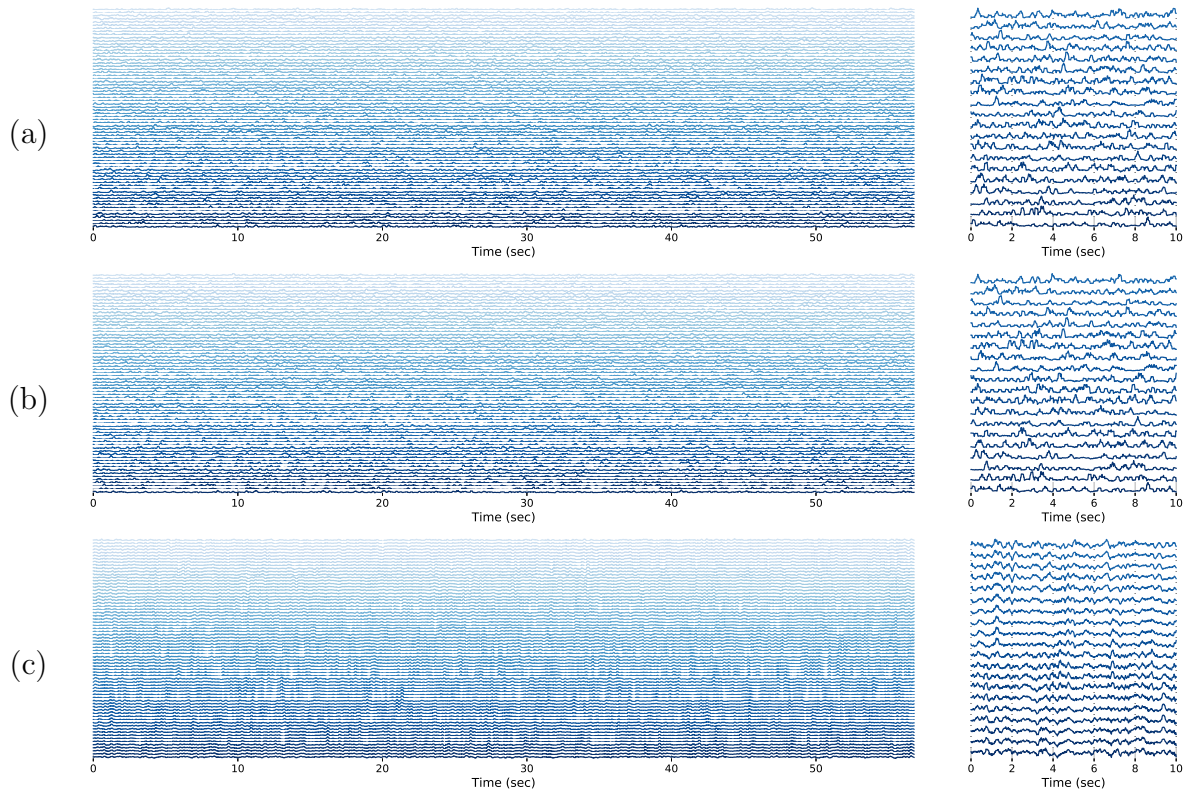


Fig. 6. Analyzing vessel #1 in Fig. 5(b). **(a)** The intensity for a single penetrating vessel in various cortical depths (sum over annotated region). Deeper layers (larger z values) are at the bottom and marked with darker colors. **(b)** The intensity after the data was multiplied by the time-collapsed segmentation mask obtained with the ACWE network.²³ **(c)** The output of our time-varying segmentation mask. **(right)** zoom-in plots of the subfigure on the left.

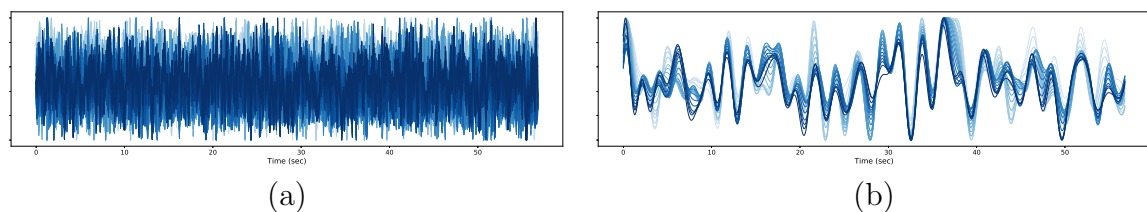


Fig. 7. The following figures refer to vessel #1 in Fig. 5(b). **(a)** The temporal sequence obtained by our method. **(b)** the result of applying a temporal low-pass filter of 1 Hz.

temporal segmentation, with its x_1 and x_2 magnification. For accurate measurements, the ratio in diameters would be exactly 2. With the skeleton layer, the average ratio is 2.04 ± 0.17 . Performing an identical 4D analysis, but without this layer, the ratio is 1.33 ± 0.40 .

Correlation between depth slices While we do not have ground truth vascular data, we can expect certain properties to hold in such data. One easily tested property is the correlation in the measured vascular activity of the same penetrating artery across adjacent axial slices.

As can be seen in Fig. 6(a), the raw measurements are very noisy and show very little correlation between adjacent axial slices. This is further illustrated in Fig. 8(a), which presents the correlation coefficients matrix between different axial slices, after removing the diagonal.

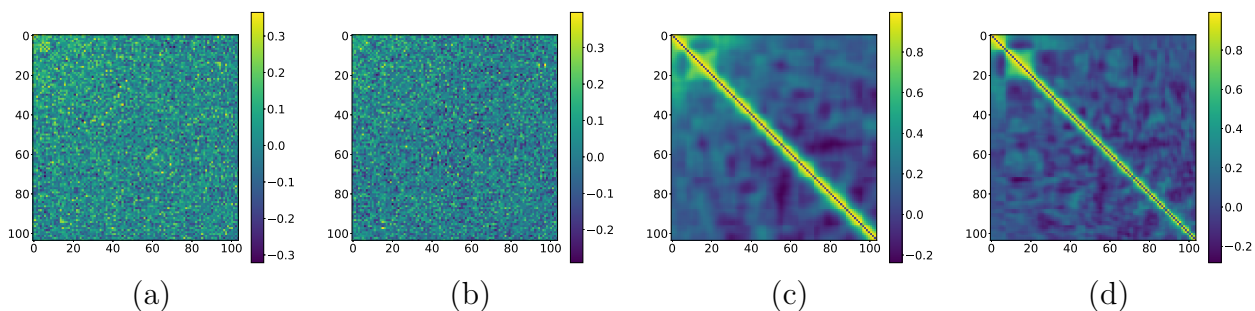


Fig. 8. Analyzing vessel #1 in Fig. 5(b). (a) The correlation between layers for the raw data. (b) Correlation for the data segmented by the ACWE network. (c) The correlation obtained by our segmentation method. (d) The correlation obtained by our method w/o the skeleton layer.

Table 1. The effect of the skeleton layer on the average Pearson correlation. z_{max} and n are the maximum z-depth and neighboring z-slices, respectively, considered for the correlation computation.

Method	z_{max} when $n = 5$				n when $z_{max} = \max$			
	25	50	75	max	1	2	5	10
W/O Skeleton layer	0.629	0.585	0.542	0.490	0.871	0.755	0.490	0.307
W/ Skeleton layer	0.743	0.731	0.728	0.721	0.956	0.907	0.721	0.485

Table 2. The increase percentage in average Pearson correlation of our method with, over our method without the skeleton layer, with respect to z_{max} and n .

z_{max} \ n	n				z_{max} \ n	n			
	1	2	5	10		1	2	5	10
25	3.8%	4.2%	6.4%	9.7%	75	18.0%	25.0%	34.4%	47.2%
50	7.6%	8.9%	13.3%	20.0%	max	30.6%	39.2%	48.2%	57.8%

The situation is not much better when multiplying the raw data by the segmentation mask of,²³ since the sparsity of the imaging modality leads to very noisy measurements (Fig. 6(b) and Fig. 8(b)). For instance, at a depth of $z = 100\mu\text{m}$ below pial surface, the brightest voxel within the brightest vessel has a time-averaged brightness of merely 0.045 photons per second, corresponding to less than 0.0004 photons per frame in time. As Fig. 6(c) shows, our method leads to a much more coherent dynamic output, which presents a large amount of correlation between adjacent axial slices. Fig. 8(c) shows that this correlation exists only between adjacent axial slices, which is a result of the gradual propagation of waves of vasodilation and vasoconstriction along the penetrating artery.

We tested our proposed skeleton layer for correlation between depth slices. In Fig. 8(c,d) we show the correlation matrix for our method with, and without, the skeleton layer, respectively. As the number of collected photons decreases with imaging depth, the resulting segmentation tends to be less coherent. This is shown in Fig. 8(d) on the bottom right side of the correlation matrix, where the correlation between neighboring slices decreases. Additionally, in Tab. 1, we show the average Pearson correlation for two cases: (i) average correlation as the maximum

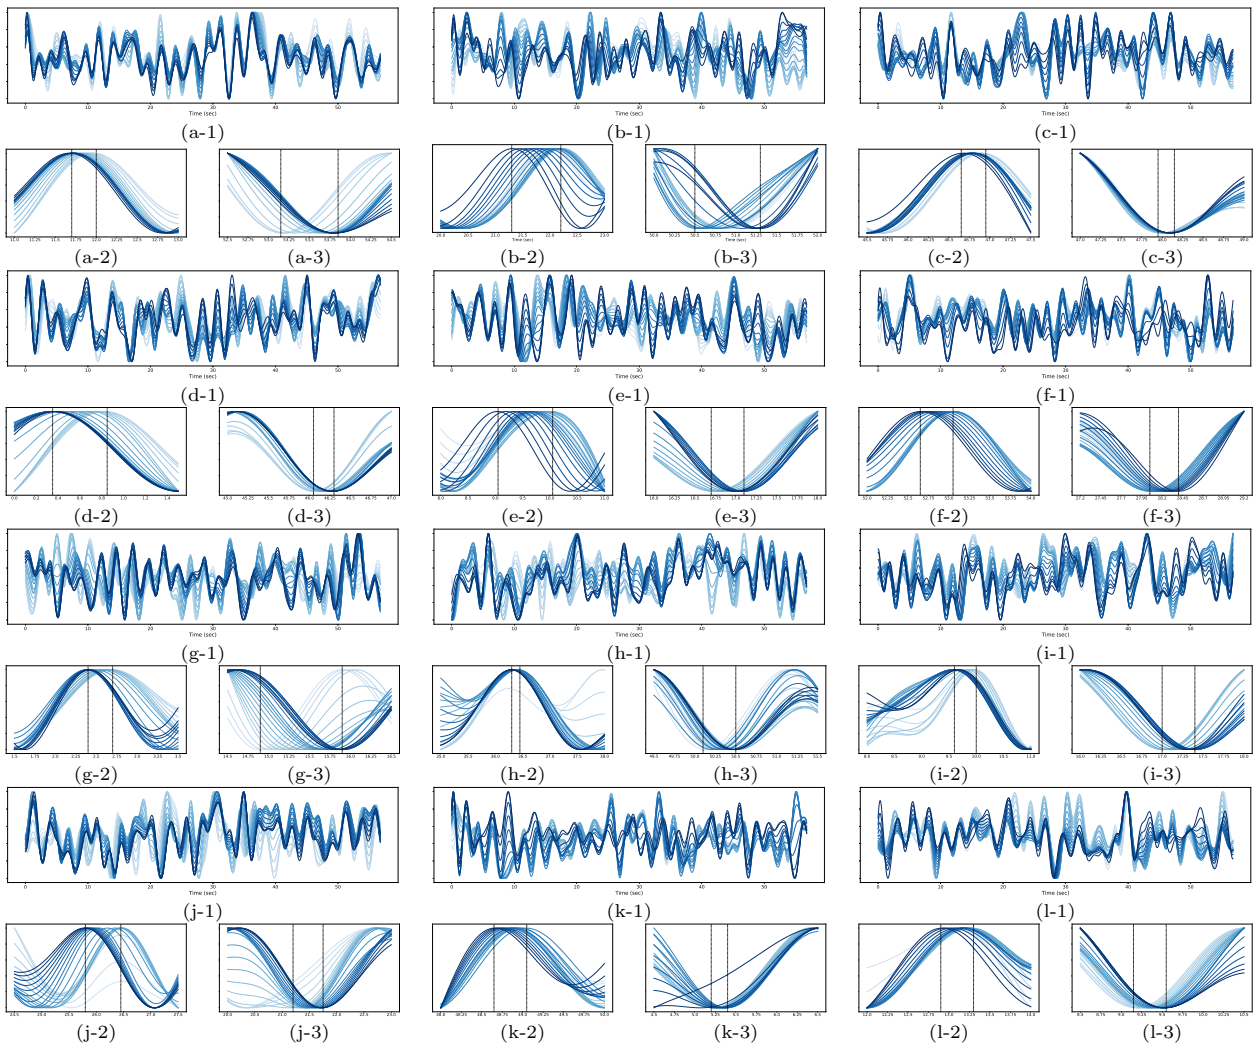


Fig. 9. The vasoconstriction and vasodilation behavior as a function of depth, extracted from vessels in Fig. 5. **(a) - (f)** and **(g) - (l)** correspond to vessels #1 to #6 in Fig. 5 (b) and (a) respectively. The measurements are taken along a single penetrating artery. In all plots, darker colors correspond to deeper cortical layers. **(*-1)** the result of applying a temporal low-pass filter with a cut-off frequency of 1 Hz. **(*-2)** zoom in on time intervals in which vasodilation begins. These results are congruent with previous observations, which were obtained using planar imaging one layer at a time and averaging over many evoked trials time-locked to an external sensory stimulus.^{7,11,12} **(*-3)** zoom in on time intervals in which vasoconstriction begins.

depth, z_{max} , increases, considering 5-neighboring slices, and (ii) the average correlation along all depth slices, considering only n neighboring slices. In Tab. 2 we show the full experiment results, showing the superiority of the skeleton layer, as an increase percentage in average Pearson correlation, especially in deeper slices and distant neighbors.

Vasodilation and vasoconstriction The dilation and constriction of the penetrating artery are depicted in Fig. 9 and Fig. 7, where we show results for the 12 annotated vessels. As can be seen in panel (a) of Fig. 7, the output of our method, when drawing multiple z-slices on the same plot, varies very fast in time. Indeed our high volumetric sampling rate, used for tracking fast neuronal activity and rigid brain motion, oversampled the propagation of vasoactivity

along the penetrating artery. We therefore applied a temporal low-pass filter with a cut-off frequency of 1 Hz. As expected, the low-pass-filtered vasoactivity traces are tightly correlated across axial depths, as shown in panel (b) of Fig. 7. Importantly, instances of vasodilation exhibit an earlier onset at deeper axial slices, as illustrated in Fig. 9(*-2). These observations of individual vasodilations tracked in several axial depths simultaneously are congruent with earlier sensory-evoked observations that were acquired and averaged one plane at a time.^{7,11,12} Conversely, instances of vasoconstriction exhibit an earlier onset at shallower axial slices, as illustrated in Fig. 9(*-3). As far as we can ascertain, this result was not observed in the past. While Fig. 9 shows only a handful of samples of vasodilation and vasoconstriction, these results are typical, and many more samples were found.

5. Conclusions

Using automated time-lapse segmentation of blood microvessels in living mouse brain we were able to track, to the best of our knowledge for the first time, how individual instances of vasodilation and vasoconstriction propagate along a penetrating artery. The observed propagation of vasodilation upwards along the penetrating artery is congruent with earlier sensory-evoked observations that were acquired one plane at a time.^{7,11,12} These propagating waves of vasoactivity along the penetrating artery are not detected by bounding the vessel with a box (Fig. 5(a)), nor by segmenting it using the time-collapsed volume (Fig. 5(b)).

Our ability to track spontaneous vasoactivity along penetrating arteries and other blood vessels in an ecologically-relevant setting paves the path towards linking it with individual action potentials. By computing the spike-triggered vasoactivity for different neuronal subtypes, we plan to distill a canonical hemodynamic response function (HRF), namely the small-signal impulse response of various classes of vessels to a single neuronal action potential.

Acknowledgment

This project has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant ERC CoG 725974). PM acknowledges the ERC (grant 639416) and the Israel Science Foundation (grant 1019/15) for support of different aspects of this project. The authors thank David Kain for conducting the mouse surgery. SG's contribution is part of a Ph.D. thesis research conducted at TAU.

References

1. P. Blinder *et al.*, The cortical angiome: an interconnected vascular network with noncolumnar patterns of blood flow, *Nature neuroscience* (2013).
2. P. S. Hosford and A. V. Gourine, What is the key mediator of the neurovascular coupling response?, *Neuroscience & Biobehavioral Reviews* (2018).
3. C. Lecrux, M. Bourourou and E. Hamel, How reliable is cerebral blood flow to map changes in neuronal activity?, *Autonomic Neuroscience* (2019).
4. C. Iadecola and R. F. Gottesman, Neurovascular and cognitive dysfunction in hypertension: epidemiology, pathobiology, and treatment, *Circulation Research* (2019).
5. A. Urban *et al.*, Understanding the neurovascular unit at multiple scales: advantages and limitations of multi-photon and functional ultrasound imaging, *Adv. drug delivery reviews* (2017).

6. K. Chhabria *et al.*, The effect of hyperglycemia on neurovascular coupling and cerebrovascular patterning in zebrafish, *J Cerebral Blood Flow & Metabolism* (2018).
7. H. Uhlirva *et al.*, Cell type specificity of neurovascular coupling in cerebral cortex, *Elife* (2016).
8. C. Echagarruga, K. Gheres and P. J. Drew, An oligarchy of no-producing interneurons controls basal and evoked blood flow in the cortex, *bioRxiv*, p. 555151 (2019).
9. P. OHerron *et al.*, Neural correlates of single-vessel haemodynamic responses in vivo, *Nature* (2016).
10. R. L. Rungta, E. Chaigneau, B.-F. Osmanski and S. Charpak, Vascular compartmentalization of functional hyperemia from the synapse to the pia, *Neuron* (2018).
11. P. Tian *et al.*, Cortical depth-specific microvascular dilation underlies laminar differences in blood oxygenation level-dependent functional mri signal, *PNAS* (2010).
12. H. Uhlirva *et al.*, The roadmap for estimation of cell-type-specific neuronal activity from non-invasive measurements, *Phil. Trans. of the Royal Society B: Biological Sciences* (2016).
13. A. P. Alivisatos, M. Chun *et al.*, The brain activity map project and the challenge of functional connectomics, *Neuron* **74** (2012).
14. N. C. Pégard, H.-Y. Liu, N. Antipa, M. Gerlock, H. Adesnik and L. Waller, Compressive light-field microscopy for 3d neural activity recording, *Optica* (2016).
15. N. Antipa, G. Kuo, R. Heckel, B. Mildenhall, E. Bostan, R. Ng and L. Waller, Diffusercam: lensless single-exposure 3d imaging, *Optica* **5**, 1 (2018).
16. A. Kazemipour *et al.*, Kilohertz frame-rate two-photon tomography, *Nature Methods* (2019).
17. Y. Wu, Y. Rivenson *et al.*, Three-dimensional propagation and time-reversal of fluorescence images, *arXiv preprint arXiv:1901.11252* (2019).
18. L. Kong, J. Tang *et al.*, Continuous volumetric imaging via an optical phase-locked ultrasound lens, *Nature methods* (2015).
19. H. Har-Gil and P. Blinder, Improving in vivo multi-photon microscopy using plug and play photon counting, *JT4A* (2019).
20. C. Stringer and M. Pachitariu, Computational processing of neural recordings from calcium imaging data, *Current opinion in neurobiology* **55**, 22 (2019).
21. Y.-R. Gao and P. J. Drew, Determination of vessel cross-sectional area by thresholding in radon space, *Journal of Cerebral Blood Flow & Metabolism* **34** (2014).
22. H. Har-Gil, L. Golgher *et al.*, Pysight: plug and play photon counting for fast continuous volumetric intravital microscopy, *Optica* **5** (2018).
23. S. Gur, L. Wolf, L. Golgher and P. Blinder, Unsupervised microvascular image segmentation using an active contours mimicking neural network, in *IEEE International Conference on Computer Vision (ICCV)*, 2019. Available online as an arXiv preprint 1908.01373.
24. P. Teikari, M. Santos *et al.*, Deep learning convolutional networks for multiphoton microscopy vasculature segmentation, *arXiv preprint arXiv:1606.02382* (2016).
25. Haft-Javaherian *et al.*, Deep convolutional neural networks for segmenting 3d in vivo multiphoton images of vasculature in alzheimer disease mouse models, *arXiv preprint 1801.00880* (2018).
26. A. P. Di Giovanna *et al.*, Whole-brain vasculature reconstruction at the single capillary level, *Scientific reports* (2018).
27. R. Damseh *et al.*, Automatic graph-based modeling of brain microvessels captured with two-photon microscopy, *j. biomedical and health informatics* (2018).
28. M. I. Todorov *et al.*, Automated analysis of whole brain vasculature using machine learning, *bioRxiv*, p. 613257 (2019).
29. T. F. Chan *et al.*, Active contours without edges, *IEEE Trans. image processing* (2001).
30. K. He *et al.*, Deep residual learning for image recognition, *CVPR* (2016).
31. C. Lantuéjoul, On the use of the geodesic metric in image analysis, *Journal of Microscopy* (2010).
32. S. Beucher, Digital skeletons in euclidean and geodesic spaces, *Signal Processing* (1994).