Supplementary Information

Bio-CM^2: Distributed computational optics for cortex-wide cellular imaging

Guorong Hu[1], Qilin Deng[1], Tianrui Qi[2], Zhixiong Chen[1], Bradley C. Rauscher[2], Nathan X. Chai[2], Daria Bogatova[2], Bethany Weinberg[3], Jesse Smith[3], Ian G. Davison[3], Anna Devor[2], Martin Thunemann[2], and Lei Tian*[1,2]

[1]Department of Electrical and Computer Engineering, Boston University, Boston, Massachusetts 02215, USA

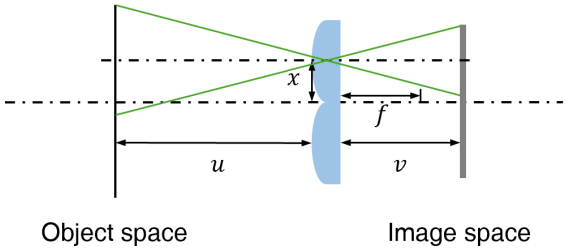
[2]Department of Biomedical Engineering, Boston University, Boston, Massachusetts, USA

[3]Department of Biology, Boston University, Boston, Massachusetts 02215, USA

*Corresponding author: leitian@bu.edu

1 Bio-CM^2 hardware

1.1 Imaging geometry design constraints



Supplementary Figure 1. Bio-CM^2 imaging geometry.

The Bio-CM^2 design procedure is illustrated in Supplementary Fig. 1. We begin by selecting a commercially available miniature CMOS sensor (4000×3000 pixels, 1.85 µm pixel size; The Imaging Source DMM 37UX226-ML), which provides both a compact form factor and sufficient pixel count. The sensor dimensions are 5.55 mm × 7.4 mm, and the shorter dimension provides the tighter geometric constraint.

Our design objective is to achieve an imaging FOV of at least 7 mm, sufficient to cover most of the dorsal cortex of a mouse, while maintaining non-overlapping image formation on the sensor. Let x denote the center position of each lens relative to the system centerline. The lens placement must simultaneously satisfy constraints on both the object plane and the sensor plane,

$$\begin{cases} 4\left(\dfrac{7}{2}-x\right)\geq 7, \\ 4\left(\dfrac{5.55}{2}-x\right)\leq 5.55, \end{cases} \tag{1}$$

where the first inequality ensures complete coverage of the target 7-mm FOV, while the second ensures that the four sub-images fit within the sensor. These constraints give

$$1.3875\ \text{mm}\leq x\leq 1.75\ \text{mm}, \tag{2}$$

which determines the allowable lens placement. Since adjacent lenses are separated by $2x$, this range also determines the maximum lens diameter that can be accommodated without physical interference or mechanical trimming of the lenses.

Next, to achieve cellular-resolution imaging with sufficient light-collection efficiency, we require an object-space numerical aperture (NA) of at least 0.15. This gives the NA constraint

$$\frac{x}{u}\geq 0.15, \tag{3}$$

where u is the object distance.

The finite-conjugate imaging geometry further satisfies

$$\begin{cases} \dfrac{\frac{5.55}{2}-x}{\frac{7}{2}-x}=\dfrac{v}{u}, \\ \dfrac{1}{u}+\dfrac{1}{v}=\dfrac{1}{f}, \end{cases} \tag{4}$$

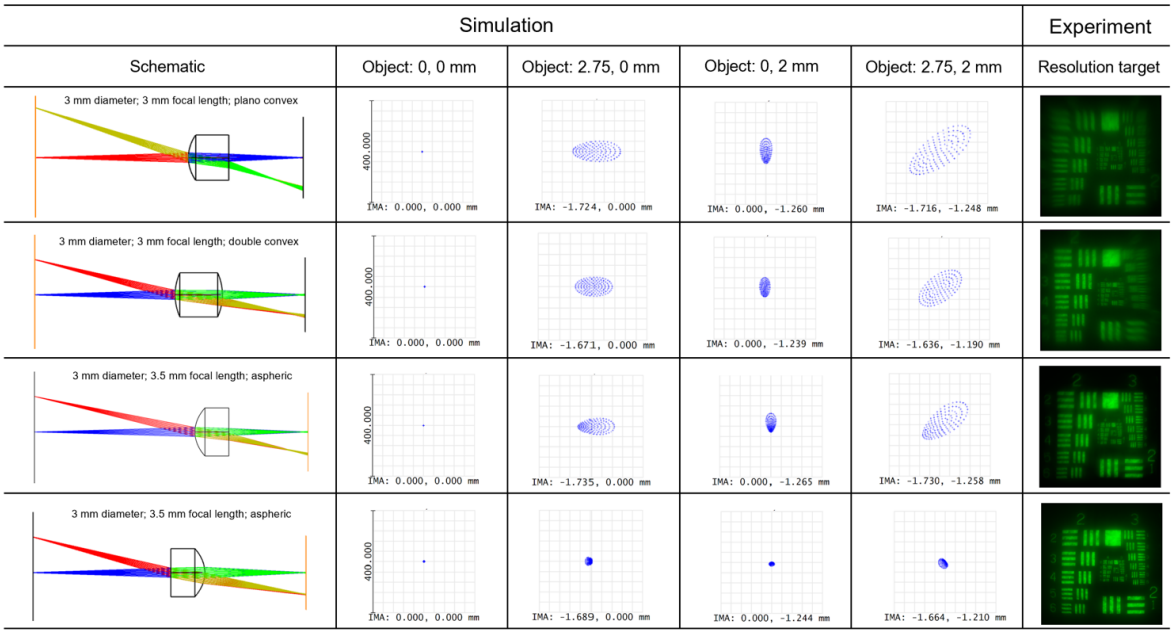
where v is the image distance and f is the focal length. The first equation sets the magnification required to map each object-space sub-FOV onto the corresponding sensor region, while the second is the thin-lens imaging equation.

Combining these relationships yields

$$f \leq \frac{20x(2.775 - x)}{18.825 - 6x}, \tag{5}$$

which defines the feasible design space relating the focal length, lens placement, and lens diameter. The resulting constraints guide the lens selection described in the next section. Once the focal length is selected, the lens position x can be determined from the geometric constraints above, after which the corresponding object distance u and image distance v are obtained from the finite-conjugate imaging equations to define the complete 2×2 Bio-CM^2 imaging geometry.

1.2 Bio-CM^2 lens selection



Supplementary Figure 2. Simulated and experimental results of selected lens candidates. The 3-mm-diameter, 3.5-mm-focal-length plano-aspheric lens, with its plano surface facing the object, provides the best overall imaging performance.

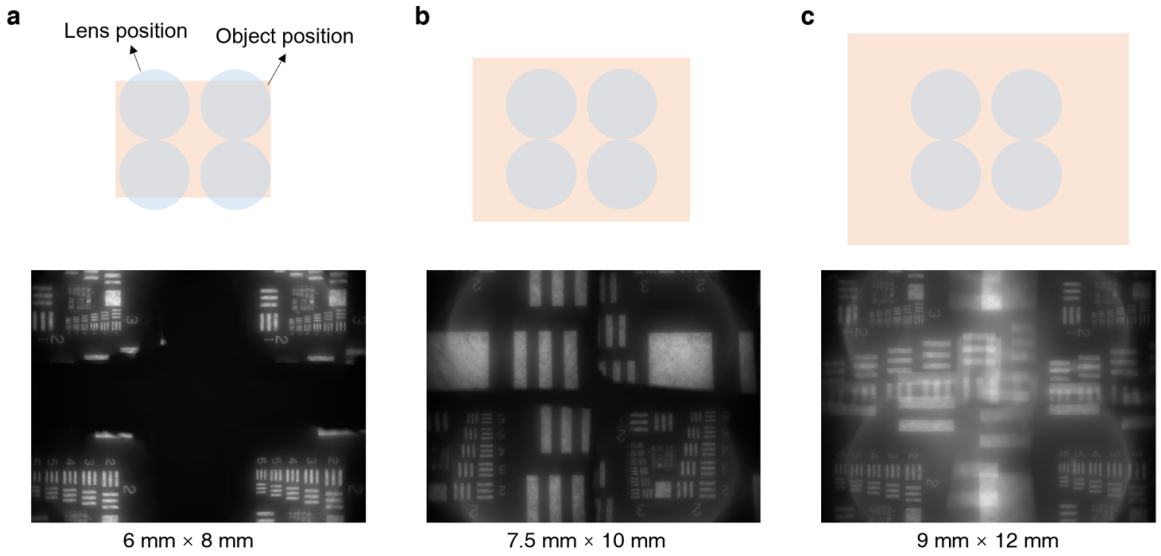
Following the design constraints above, we identified three commercially available candidate lenses, all with a 3-mm diameter: a plano-convex lens with a 3-mm focal length, a double-convex lens with a 3-mm focal length, and a plano-aspheric lens with a 3.5-mm focal length. Assuming a single-lens sub-FOV of 4×5.55 mm^2, we evaluated the three candidates using Zemax simulations. Among them, the plano-aspheric lens produced the most compact simulated spot diagrams (PSFs) for input across the local FOV and was therefore selected for further optimization.

Therefore, our final design uses the following parameters: f = 3.52mm, u = 9.02 mm, v = 5.77 mm, lens diameter = 3mm, horizental distance x = 2mm, vertical distance x = 1.5mm.

We next compared the two orientations of the plano-aspheric lens. Orienting the plano surface toward the object slightly degrades the on-axis PSF but noticeably improves the off-axis PSFs, resulting in better imaging performance across the full single-lens sub-FOV. We therefore selected the plano-aspheric lens with its plano surface facing the object for the final Bio-CM^2 design.

Finally, we experimentally validated the Zemax predictions using a fluorescence resolution target imaged with a single-lens system having the same imaging geometry and magnification. The experimental images closely reproduce the trends predicted by the simulations, confirming both the lens choice and the selected orientation (Supplementary Fig. 2).

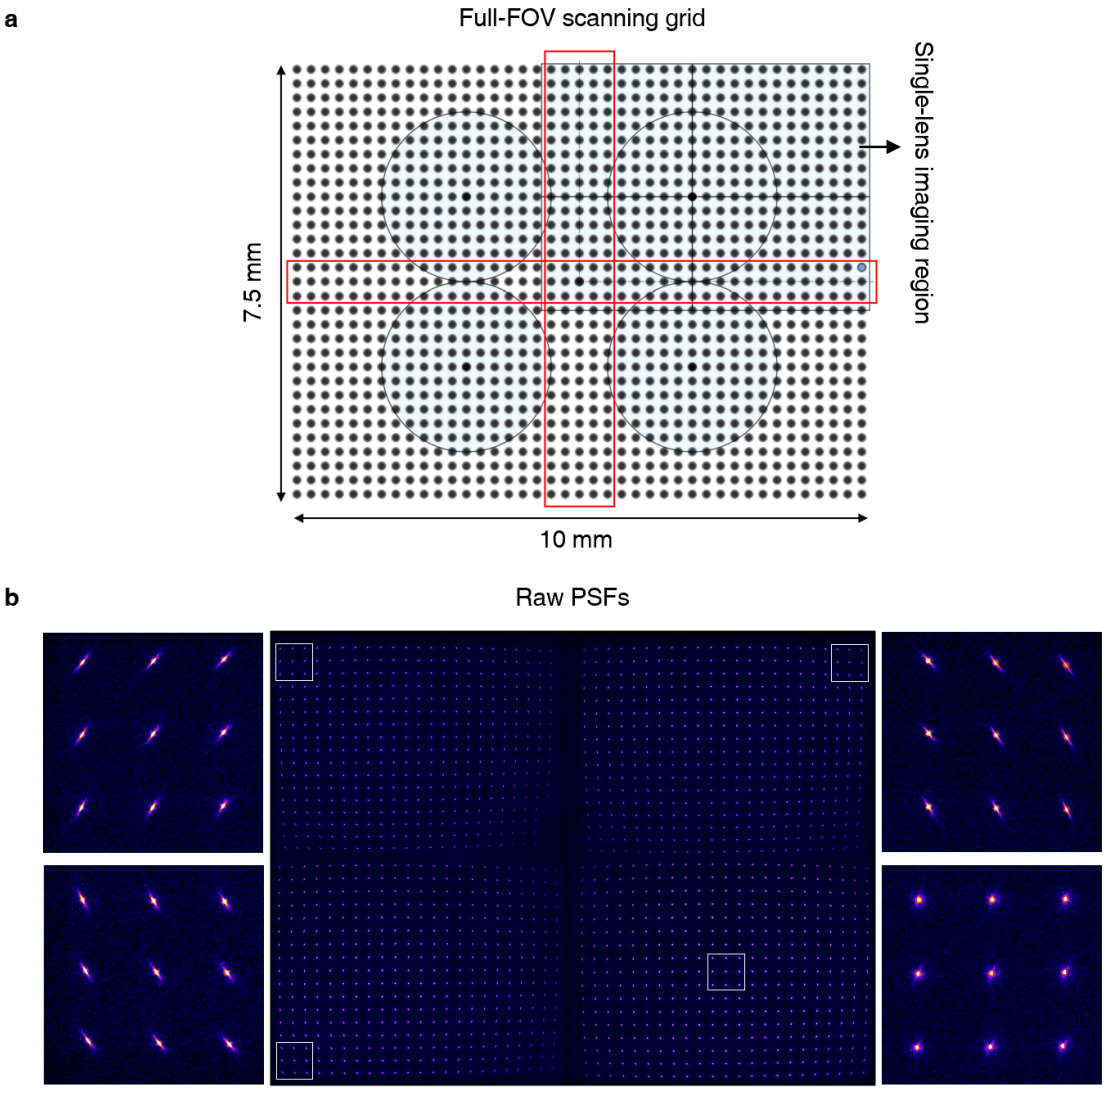
1.3 Experimental FOV characterization



Supplementary Figure 3. Experimental validation of the Bio-CM^2 imaging geometry. The designed imaging geometry achieves a 7.5×10 mm^2 field of view while maintaining non-overlapping measurements from the four imaging channels. (a) For objects smaller than the designed field of view, the four sub-images remain spatially separated on the sensor. (b) At the designed field of view, the four sub-images are tiled without overlap, corresponding to the nominal operating condition. (c) When the object exceeds the designed field of view, adjacent sub-images begin to overlap, marking the designed imaging limit.

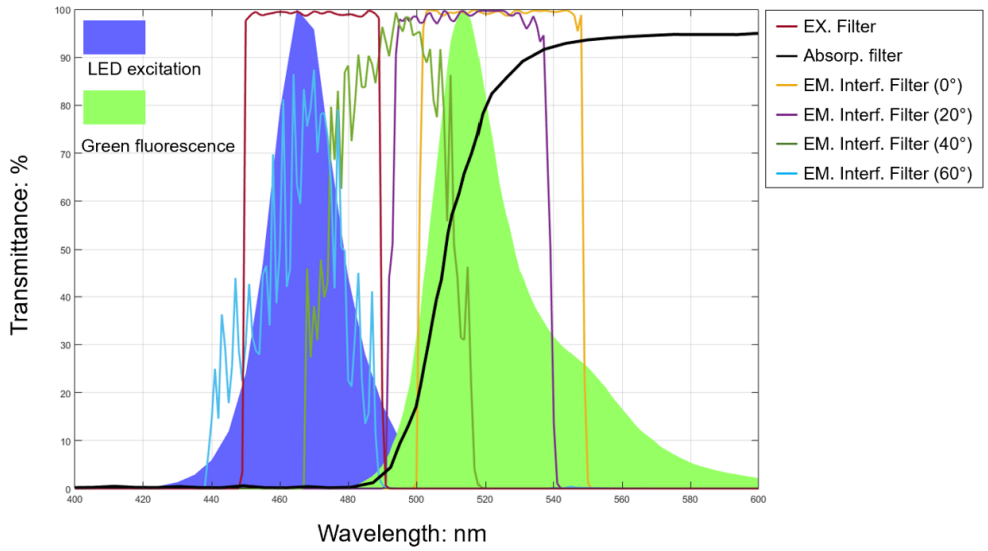
The designed FOV was experimentally validated using a fluorescence resolution target masked to different effective imaging areas. As predicted by the geometric analysis, the four imaging channels remain completely separated on the sensor within the designed 7.5×10 mm^2 FOV (Supplementary Fig. 3a,b). When the object size exceeds the design limit, adjacent sub-images begin to overlap on the sensor (Supplementary Fig. 3c). These results experimentally verify both the distributed imaging geometry and the derived design constraints.

1.4 PSF calibration



Supplementary Figure 4. PSF calibration. (a) Full-field PSF calibration using a fluorescent bead scanned over the entire 7.5 mm × 10 mm FOV. The blue shaded region denotes the nominal imaging area of a single lens. The red horizontal and vertical bands represent overlap regions where object points are simultaneously captured by two adjacent lens channels; at the center of the FOV, each point is captured by all four lens channels. The complete scan consists of a 31 × 41 grid with a 250 µm step size. (b) Representative raw PSFs acquired across the full FOV. The center panel shows all measured PSFs, and the enlarged views highlight representative PSFs from different imaging locations, illustrating the spatially varying optical response used for system calibration.

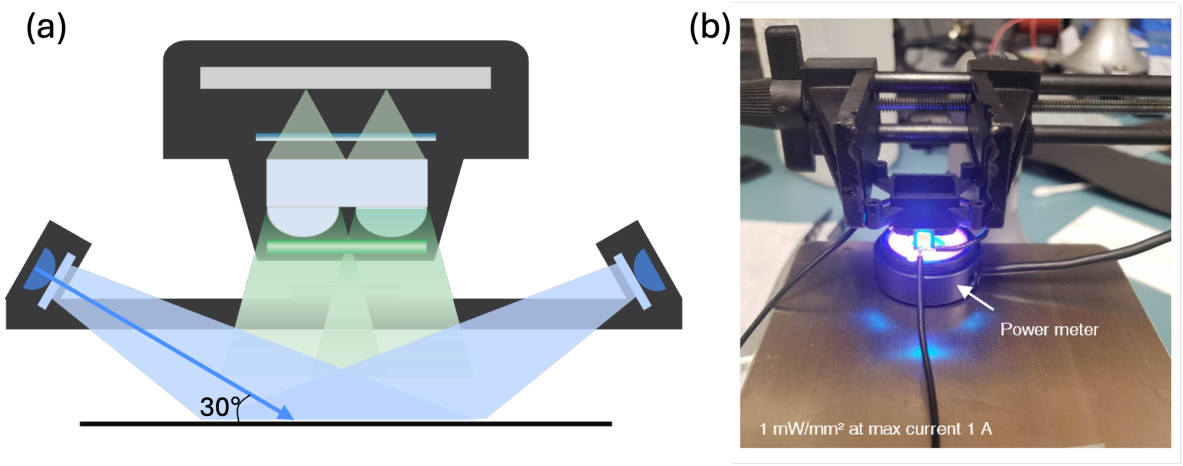
1.5 Bio-CM^2 fluorescence filter design



Supplementary Figure 5. Spectral properties of Bio-CM^2 filters. Transmission spectra of the interference-based emission filter at incidence angles of 0°, 20°, 40°, and 60°, illustrating its angle-dependent spectral shift.

The fluorescence filter design follows our previous work [1, 2]. On the illumination path, an interference-based excitation filter is placed immediately after the LED to spectrally purify the excitation light. On the imaging path, however, fluorescence rays are collected over a wide range of incidence angles. Consequently, the passband of the interference-based emission filter shifts toward shorter wavelengths at oblique incidence (Supplementary Fig. 5), reducing excitation rejection and allowing stray excitation light to leak into the measurements. To suppress this angle-dependent leakage, an incidence-independent long-pass absorption filter is placed after the interference filter, substantially reducing the background signal.

1.6 Bio-CM^2 illumination power measuring setup



Supplementary Figure 6. Bio-CM^2 illumination geometry and irradiance measurement setup. (a) Schematic of the illumination and imaging geometry. Four symmetrically arranged oblique excitation arms (two shown for clarity) deliver blue excitation light to the sample at an elevation angle of approximately 30° relative to the imaging plane, while the emitted fluorescence (green) is collected through the central imaging optics. (b) Photograph of the irradiance measurement setup, with a power meter positioned at the Bio-CM^2 imaging plane. At the maximum LED drive current of 1 A, the measured irradiance is approximately 1 mW/mm^2, sufficient for cortex-wide one-photon in-vivo neural imaging [3].

1.7 Bio-CM^2 costs

Supplementary Table 1. Bill of materials for the Bio-CM^2 prototype (Retrieved on 06/23/2026).

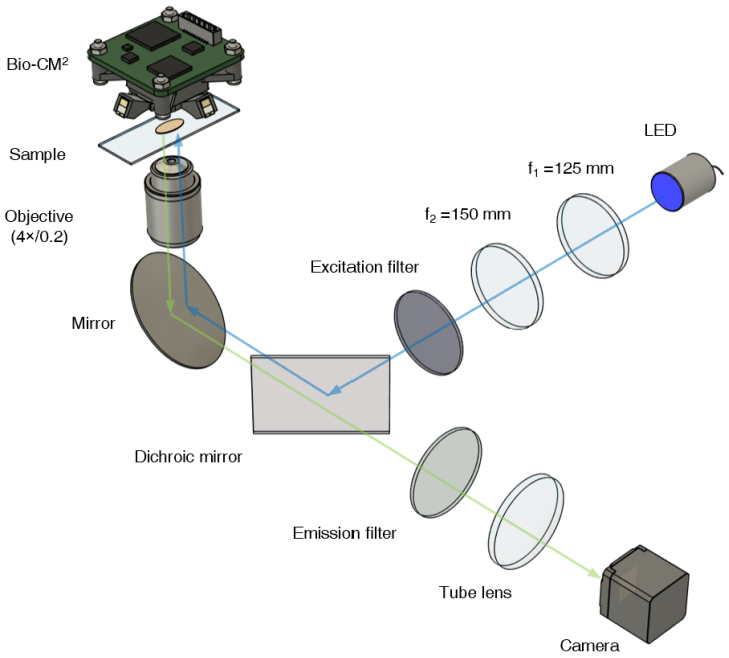
Item	Part Number	Vendor	Price	Quantity	Cost
Lenses	353515	Edmund Optics	$89	4	$356
Camera	DMM 37UX226-ML	The Imaging Source	$328	1	$328
Emission filter pair	ET525-50m	Chroma	$182.5	1	$182.5
	54-467	Edmund Optics	$131	1	$131
Excitation filter	ET470/40x	Chroma	∼$40	4	$160
LED	LXML-PB01-0040	Lumileds	∼$3	4	$12
				Total Cost:	**$1169.5**

1.8 Comparison of FOV and spatial resolution between Bio-CM^2 and other miniature imaging devices

Supplementary Table 2. Comparison of representative miniature fluorescence imaging systems.

System	FOV (mm^2)	Lateral Resolution (μm)
mini-mScope [4]	8×10	39-56
mini-MCAM [5]	$4 \times 4.5 \times 2.55$	9.9
CortexCAM [6]	48	13.9
SOMM [7]	3.6×3.6	4
Miniscope-LFOV [8]	3.6×2.7	3.5-6.2
MiniXL [9]	3.5×3.5	4.4
Bio-FlatScope [10]	4×4	9
Bio-FlatScopeNHP [11]	4.5×4.5	10
Kiloscope [12]	4.8×3.6	5
DeepInMiniscope [13]	4×6	10
CM^2 [14]	8.1×7.3	7
CM^2 V2 [1]	6.7×6.7	6
Bio-CM^2	7.5×10	6

1.9 Benchtop system



Supplementary Figure 7. Benchtop calibration and validation system. The setup comprises a shared epi-illumination path and two imaging paths sharing a common sample plane. A reference fluorescence microscope images the sample from below, while Bio-CM^2 images the same specimen from above. The dual-imaging configuration enables precise optical alignment, PSF calibration, quantitative performance characterization, and direct comparison of biological imaging results.

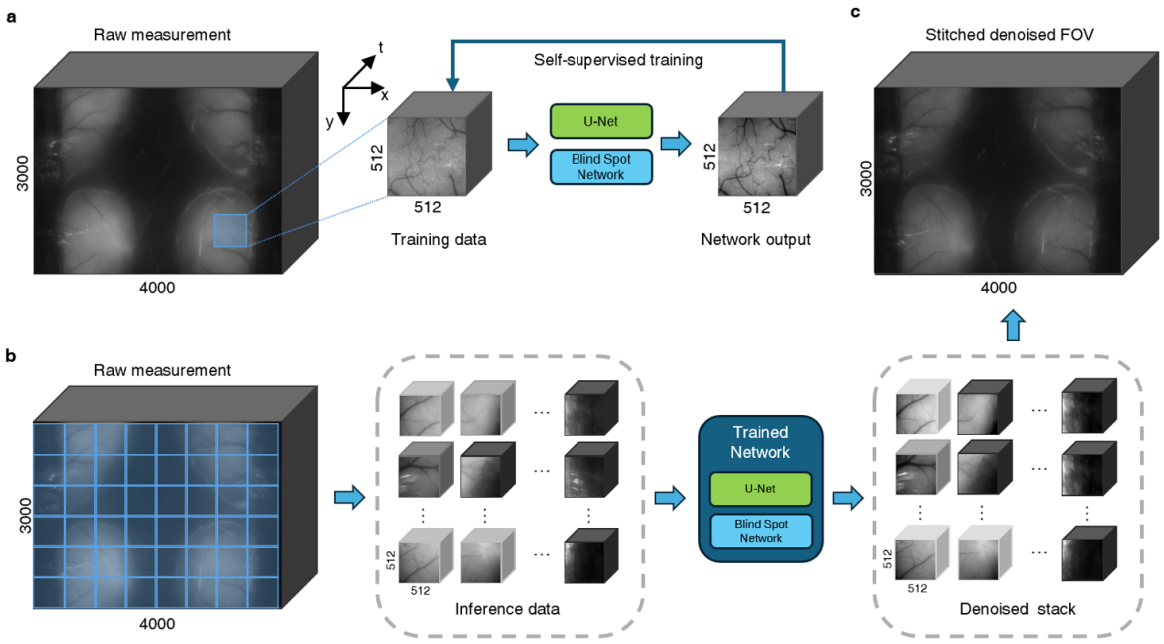
2 Additional details on algorithms

2.1 SUPPORT denoising for in vivo neural imaging data

Our in vivo neural imaging dataset is large in both temporal and spatial dimensions ($1798(t) \times 4000(x) \times 3000(y)$), making direct denoising of the full volume infeasible on available GPUs. We therefore crop the video along the spatial dimensions into 48 patches, each of size $1798(t) \times 512(x) \times 512(y)$. A 12-pixel overlap between adjacent patches is used for alpha blending during stitching to mitigate boundary effects.

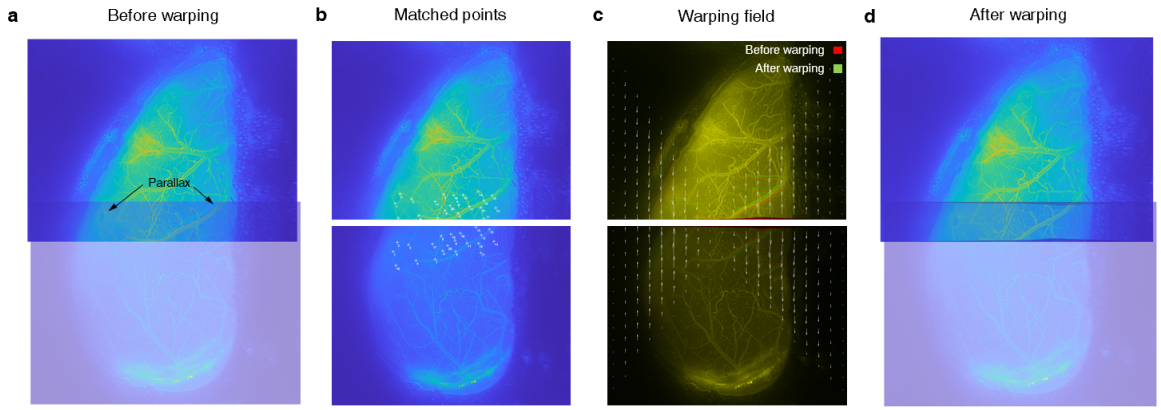
Empirically, the lower-right lens-center patch shows the fewest reflective artifacts, so we use it for training and generalize the resulting model to all remaining patches.

The SUPPORT network [15] is trained in PyTorch (2.4.0) with CUDA (12.5) on an NVIDIA RTX 4090 GPU and an AMD Ryzen Threadripper PRO 5955WX 16-core CPU. Starting with the selected $1798(t) \times 512(x) \times 512(y)$ video, we first normalize it by subtracting its mean and dividing by its standard deviation. From the normalized data, we generate training samples of size $11(t) \times 128(x) \times 128(y)$, where adjacent samples overlap by $64(x) \times 64(y)$ spatially and each sample consists of a target frame together with its 5 preceding and 5 succeeding frames temporally. To ensure complete temporal support, the first and last 5 frames are excluded from target-frame selection. These factors result in 1788 temporal positions and 7 patch positions along each spatial dimension, yielding a total of 87,612 training samples. With a batch size of 16, each epoch contains 5476 iterations. During training, each sample is randomly augmented using vertical flips, rotations by multiples of 90°, or a combination of both. The network is optimized using the Adam optimizer with a learning rate of 5×10^{-4} and no weight decay. Random seeds are set to 0 across all relevant libraries to ensure reproducibility. The model is trained for 10 epochs. Training on one patch of size $1798(t) \times 512(x) \times 512(y)$ takes 0.69 hours, and inference per patch takes 93 seconds. Zero padding is used during both training and inference to preserve input–output size consistency.



Supplementary Figure 8. SUPPORT denoising pipeline for Bio-CM^2. (a) A patch centered on the lower-right lens is extracted from the raw video for self-supervised training using a U-Net combined with a blind-spot network. (b) After training, the raw video is partitioned into spatially overlapping patches, which are processed by the trained network to generate denoised stacks. (c) The denoised stacks are stitched using their overlapping regions to reconstruct the denoised full FOV.

2.2 Parallax compensation for image stitching



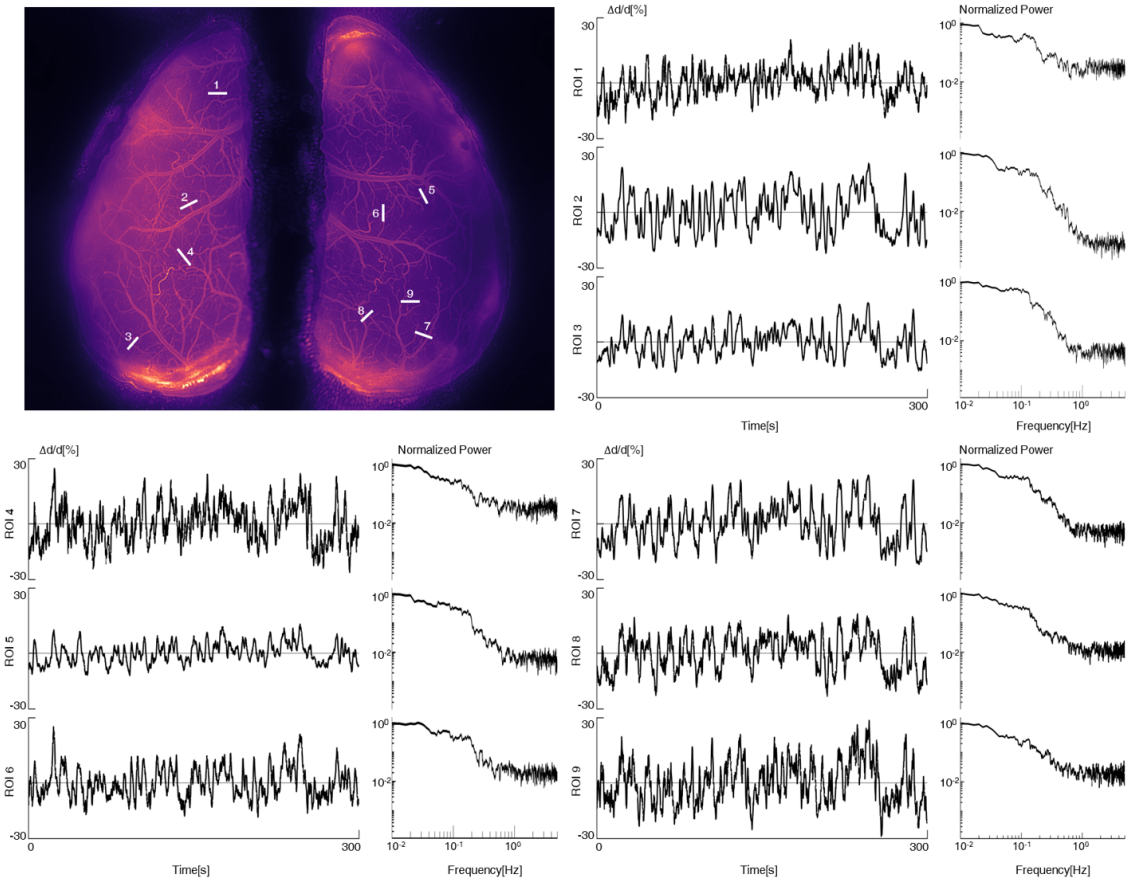
Supplementary Figure 9. Parallax compensation prior to image stitching. (a) Representative overlapping *in vivo* vascular image pair exhibiting parallax-induced misalignment between adjacent views. (b) Manually selected corresponding vessel feature points used to estimate the translational warp fields. (c) Pixel-wise translational warp fields generated by linear interpolation. Arrows indicate the local displacement vectors, and the overlay compares the image pair before (red) and after (green) warping. (d) Overlapping image pair after parallax compensation, showing improved alignment before alpha blending.

For quasi-2D samples, the four sub-images can be stitched directly by alpha blending because corresponding structures are well aligned within the overlap regions. However, *in vivo* samples such as the curved mouse cortex exhibit appreciable depth variation, causing parallax between adjacent imaging channels. Consequently, corresponding structures are displaced within the overlap regions and cannot be blended seamlessly without additional correction.

To compensate for this effect, corresponding feature points are manually selected from the overlap regions of neighboring sub-images. Blood vessels provide high-contrast anatomical landmarks that are readily identifiable across adjacent views, enabling reliable point correspondence. For each pair of neighboring sub-images, the displacement vectors between matched feature points are linearly interpolated to generate dense pixel-wise translational warp fields. Each sub-image is then spatially warped according to the interpolated displacement field using bilinear interpolation, before alpha blending.

Supplementary Fig. 9 illustrates the procedure using representative *in vivo* vascular images. Before warping, adjacent views exhibit noticeable parallax-induced misalignment (Supplementary Fig. 9a). Manually selected vessel correspondences (Supplementary Fig. 9b) are used to estimate the translational warp fields (Supplementary Fig. 9c), resulting in substantially improved alignment after warping (Supplementary Fig. 9d). Once parallax is compensated, the corrected sub-images can be stitched seamlessly by alpha blending.

3 Additional vessel ROI analysis



Supplementary Figure 10. Additional vessel-diameter measurements across the dorsal cortex. Representative vessel ROIs distributed across the dorsal cortical vasculature are labeled in the upper left. For each ROI, the corresponding relative vessel diameter change ($\Delta d/d$) and normalized power spectrum are shown. The measurements exhibit consistent vasomotion dynamics across vessels distributed throughout the dorsal cortex, further demonstrating that Bio-CM^2 enables reliable quantification of single-vessel dynamics at cortex-wide scale.

4 Additional details on CNMF-E

4.1 CNMF-E pipeline

The neuronal extraction pipeline is summarized in Supplementary Fig. 11. Specifically, CNMF-E [16], implemented in CaImAn v1.13.1 [17], models an input fluorescence movie $\mathbf{Y} \in \mathbb{R}^{d \times T}$ (Supplementary Fig. 11a) as

$$\mathbf{Y} = \mathbf{A}\mathbf{C}^\top + \mathbf{B}^c + \mathbf{B}^f + \mathbf{E}$$

where d is the number of pixels and T is the number of frames. $\mathbf{B}^c = \mathbf{b}_0\mathbf{1}^\top$ is rank-one constant baseline where $\mathbf{b}_0 \in \mathbb{R}^d$ (Supplementary Fig. 11e) and $\mathbf{1} \in \mathbb{R}^T$. $\mathbf{B}^f \in \mathbb{R}^{d \times T}$ (Supplementary Fig. 11f) represents the fluctuating background activity, modeled by expressing pixel i at frame t background as a linear combination of the background fluorescence from pixels within its ring neighborhood at frame t,

$$b^f_{i,t} = \sum_{j \in \Omega_i} w_{i,j} b^f_{j,t} \qquad \Omega_i = \left\{ j \mid \|\mathbf{p}_i - \mathbf{p}_j\|_2 \in [l_n, l_n + 1] \right\}$$

where $\mathbf{p}_i \in \mathbb{R}^2$ denotes the two-dimensional spatial coordinate of pixel i, $w_{i,j}$ is the fitted weight for neighboring pixel j, and Ω_i is the ring neighborhood of pixel i. The resulting weight matrix $\mathbf{W} \in \mathbb{R}^{d \times d}$ is sparse, with $w_{i,j} = 0$ for $j \notin \Omega_i$. We used $l_n = 10$ pixels (28.9 µm), corresponding to a one-pixel-wide discrete ring at an approximately 10–11-pixel distance from pixel i. In practice, CNMF-E estimates the fluctuating background from the residual after subtracting the neuronal signal and constant baseline:

$$\mathbf{B}^f = \mathbf{W}\left(\mathbf{Y} - \mathbf{A}\mathbf{C}^\top - \mathbf{B}^c\right).$$

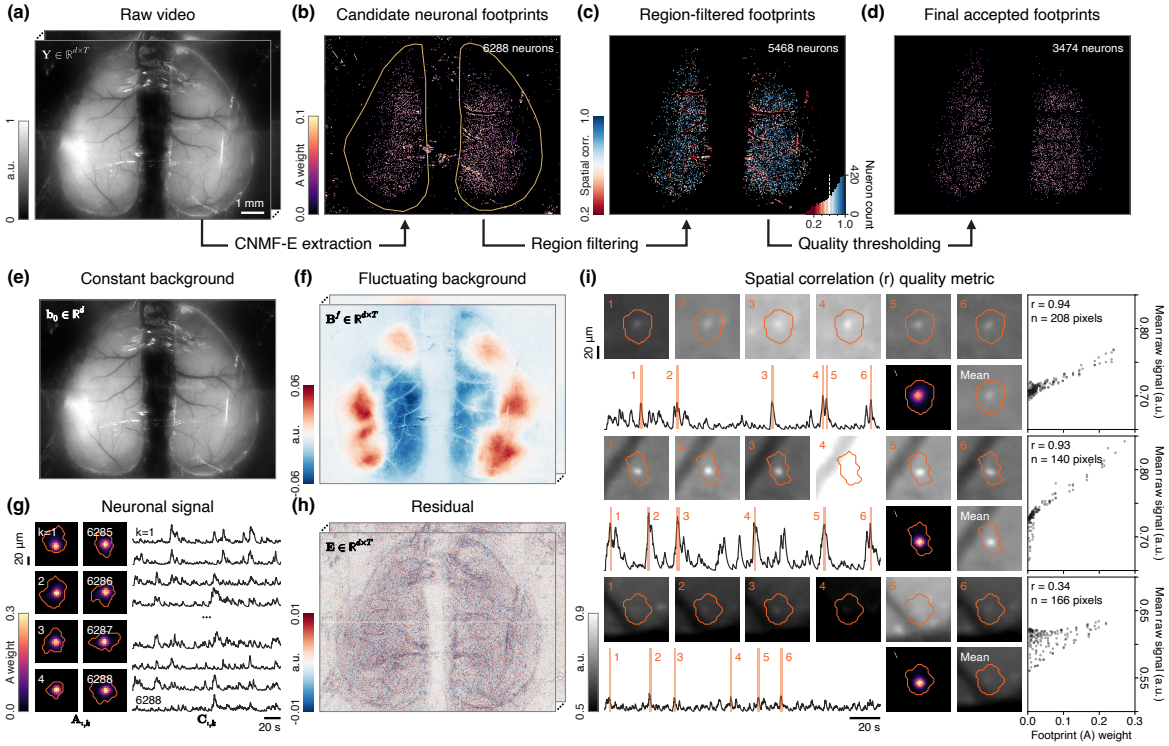
The neuronal signal is represented by $\mathbf{A}\mathbf{C}^\top$, where $\mathbf{A} = [\mathbf{a}_1, \ldots, \mathbf{a}_K] \in \mathbb{R}_+^{d \times K}$ contains the non-negative, sparse, and spatially localized neuronal footprints, and $\mathbf{C} = [\mathbf{c}_1, \ldots, \mathbf{c}_K] \in \mathbb{R}_+^{T \times K}$ contains their corresponding temporal calcium traces (Supplementary Fig. 11g). To impose calcium-dynamics constraints on the temporal traces, each $\mathbf{c}_k \in \mathbb{R}_+^T$ is modeled using a stable second-order autoregressive (AR(2)) process:

$$c_{k,t} = s_{k,t} + \gamma_{k,1} c_{k,t-1} + \gamma_{k,2} c_{k,t-2},$$

where $s_{k,t} \geq 0$ is the inferred sparse spike/event drive at frame t, and $\gamma_{k,1}$ and $\gamma_{k,2}$ are component-specific AR coefficients estimated during temporal fitting. The AR(2) model allows the fitted calcium transient to capture both a finite rise and decay following a spike. For each retained component, the calcium response to a single action potential (AP) is estimated by applying a unit impulse to the fitted AR(2) model and evaluating the resulting temporal response. After normalization to unit peak amplitude, t_{peak} is defined as the time from the impulse to the response maximum, and $t_{1/2}$ is defined as the post-peak time required for the response to decay to half of its peak amplitude.

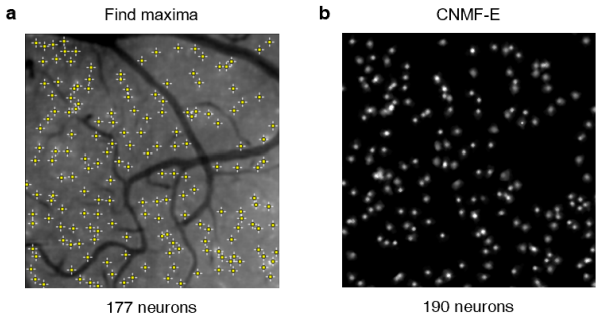
Candidate components are first filtered using manually defined cortical masks, retaining only those located within the delineated cortical regions; this step reduces the initial 6,288 CNMF-E components to 5,468 region-filtered components (Supplementary Fig. 11b,c). We then assign each region-filtered component a spatial-correlation score r_k to quantify how well its inferred footprint is supported by the raw movie. For component k, active frames are identified from its fitted temporal trace $\mathbf{c}_k$, and the raw fluorescence over these frames is averaged to form an activity-associated fluorescence image. The score r_k is computed as the Pearson correlation between the inferred footprint weights and this mean raw-fluorescence pattern over pixels within the footprint. Thus, components with high r_k have larger footprint weights at pixels that are also brighter in the raw movie when the component is

active. In practice, components with low r_k are typically false-positive detections near blood vessels, where *in vivo* motion and hemodynamic artifacts often lead to erroneous source identification. The region-filtered components are displayed in Supplementary Fig. 11c with each footprint colored by r_k, and components with $r_k \geq 0.60$ are retained as final accepted neurons, yielding 3,474 components (Supplementary Fig. 11d). Representative examples in Supplementary Fig. 11i show two accepted components with high spatial-correlation scores ($r = 0.94$ and $r = 0.93$) and one rejected component with poor footprint–fluorescence agreement ($r = 0.34$).



Supplementary Figure 11. CNMF-E-based neuronal signal extraction and quality control. **a**, Representative frame from the reconstructed full-FOV calcium-imaging movie used as input to CNMF-E. **b**, Candidate neuronal footprints identified by CNMF-E across the full FOV ($n = 6,288$ components); yellow outlines indicate the manually delineated cortical regions used for region filtering. **c**, Region-filtered components ($n = 5,468$) after retaining components located within the cortical masks, with each footprint colored by its spatial-correlation score r. **d**, Final accepted neuronal footprints after applying the spatial-correlation quality threshold ($r \geq 0.60$; $n = 3,474$ components). **e**–**h**, Example CNMF-E decomposition of the fluorescence movie into constant background ($\mathbf{B}^c$; e), fluctuating background ($\mathbf{B}^f$; f), neuronal signal ($\mathbf{A}, \mathbf{C}$; g), and residual ($\mathbf{E}$; h). **i**, Spatial-correlation quality metric for representative components. Active periods in each fitted temporal trace are used to compute an activity-associated mean raw-fluorescence image, which is compared with the inferred footprint over the footprint support. Two accepted components show strong footprint–fluorescence agreement ($r = 0.94$ and $r = 0.93$), whereas one rejected component shows weak agreement ($r = 0.34$). The reported n values indicate the number of pixels in each footprint support.

4.2 Independent validation of CNMF-E neuron counts using local maxima

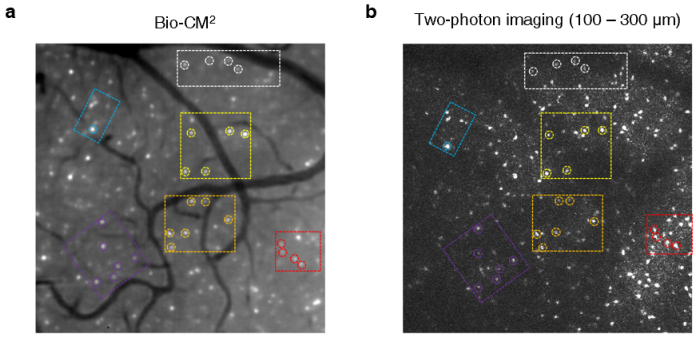


Supplementary Figure 12. Independent validation of CNMF-E neuron counts using local-maxima detection. (a) Neuron centers estimated by a simple local-maxima detection method from the enlarged view outlined by the blue dashed box in Fig. 6a. This approach provides an independent estimate of the number of resolvable neurons without model-based source extraction, yielding 177 detected neurons. (b) Spatial footprints of 190 neurons extracted by CNMF-E from the same region. The close agreement between these two independent estimates supports the parameter selection and neuron counts obtained by CNMF-E.

To independently assess the neuron counts obtained by CNMF-E [16, 17], we applied a simple local-maxima detection method to the same cortical region (Supplementary Fig. 12). Local maxima were identified from the temporal standard deviation projection of the reconstructed video, providing a model-free estimate of the number of resolvable neurons. This approach detected 177 neurons, in close agreement with the 190 neurons identified by CNMF-E, supporting the robustness of the CNMF-E parameter selection.

The local-maxima method, however, is intended only as a validation tool and has important limitations. Because it operates on a temporally projected image, it discards the rich temporal dynamics contained in the calcium imaging video and cannot separate neurons that overlap spatially but exhibit distinct temporal activity. In contrast, CNMF-E jointly exploits the spatial and temporal structure of the full video to separate fluctuating background fluorescence from neuronal signals, resolve overlapping neurons, and recover individual neuronal activity traces.

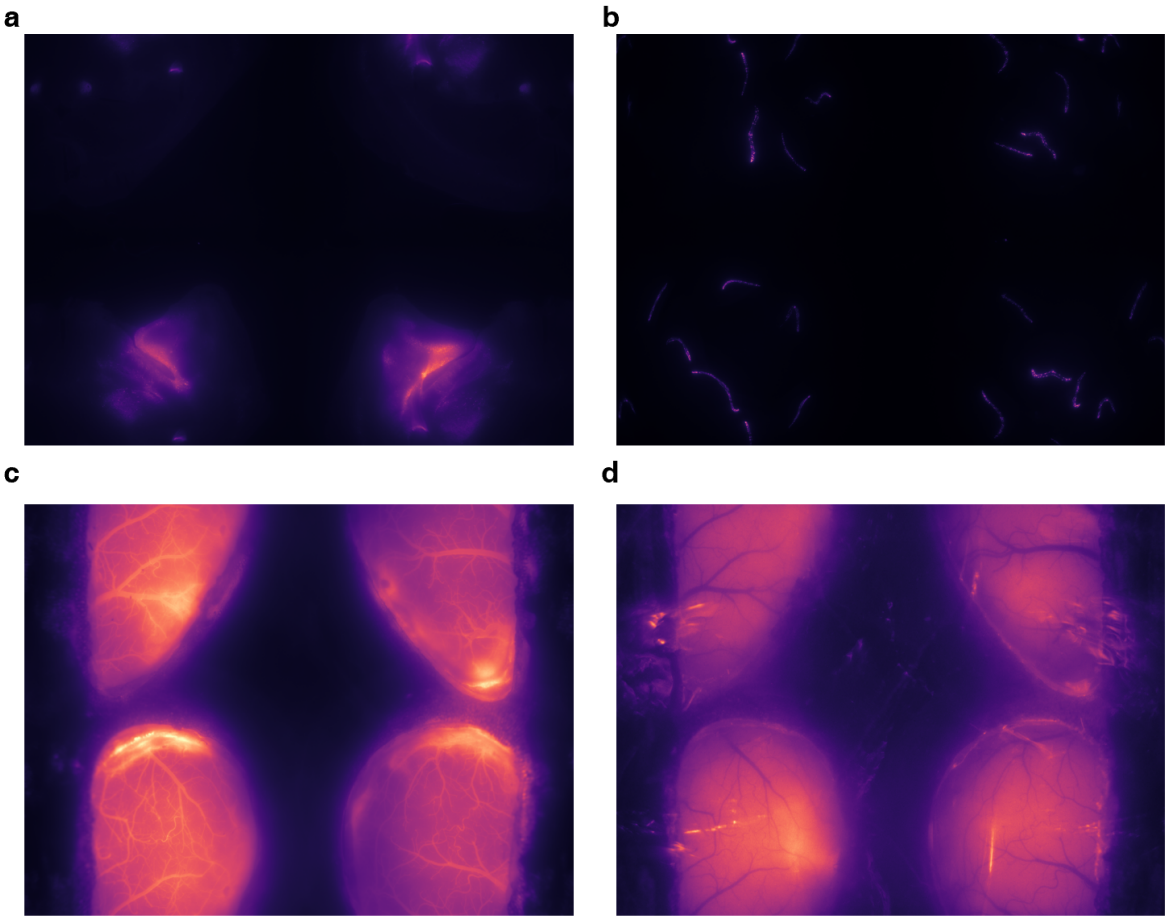
5 Cross-validation of Bio-CM^2 neuron detection with two-photon microscopy



Supplementary Figure 13. Single-neuron cross-validation between Bio-CM^2 and two-photon microscopy. (a) Enlarged Bio-CM^2 view corresponding to the blue dashed box in Fig. 6a. The image is displayed as the temporal standard deviation projection of the calcium imaging video. (b) Maximum-intensity projection of the corresponding two-photon image stack over a depth range of 100–300 µm. Although the Bio-CM^2 and two-photon datasets were acquired independently and visualized using different image representations (temporal standard deviation versus maximum-intensity projection), individual neurons exhibit strong spatial correspondence between the two modalities. Colored boxes and circled neurons are shown only to facilitate visual comparison; additional corresponding neurons can also be identified.

To validate neuron detection by Bio-CM^2 at the single-cell level, we compared the reconstructed calcium image with a corresponding *in vivo* two-photon imaging dataset acquired from the same cortical region. Because Bio-CM^2 primarily captures neurons within cortical layer 2/3 under sparse labeling conditions, the two-photon image was represented by a maximum-intensity projection over a depth range of 100–300 µm, which provided the closest correspondence. As shown in Supplementary Fig. 13, individual neurons identified by Bio-CM^2 exhibit strong spatial agreement with the two-photon reference, supporting the accuracy of single-neuron detection.

6 Raw measurements in main Figure 3–6



Supplementary Figure 14. Representative raw measurements corresponding to main Figs. 3–6. (a) Raw measurement of the fixed mouse brain slice. (b) Raw measurement of the *C. elegans*. (c) Raw measurement from the cortex-wide vascular imaging. (d) Raw measurement from the cortex-wide neuronal calcium imaging.

References

[1] Yujia Xue, Qianwan Yang, Guorong Hu, Kehan Guo, and Lei Tian. Deep-learning-augmented computational miniature mesoscope. *Optica*, 9(9):1009–1021, 2022.

[2] Qianwan Yang, Ruipeng Guo, Guorong Hu, Yujia Xue, Yunzhe Li, and Lei Tian. Wide-field, high-resolution reconstruction in computational multi-aperture miniscope using a fourier neural network. *Optica*, 11(6):860–871, 2024.

[3] Isaac V. Kauvar, Timothy A. Machado, Elle Yuen, John Kochalka, Minseung Choi, William E. Allen, Gordon Wetzstein, and Karl Deisseroth. Cortical observation by synchronous multifocal optical sampling reveals widespread population encoding of actions. *Neuron*, 107(2):351–367.e19, 2020.

[4] Mathew L Rynes, Daniel A Surinach, Samantha Linn, Michael Laroque, Vijay Rajendran, Judith Dominguez, Orestes Hadjistamoulou, Zahra S Navabi, Leila Ghanbari, Gregory W Johnson, Mojtaba Nazari, Majid H Mohajerani, and Suhasa B Kodandaramaiah. Miniaturized head-mounted microscope for whole-cortex mesoscale imaging in freely behaving mice. *Nature Methods*, 18(4):417–425, 2021.

[5] Jia Hu, Arun Cherkkil, Daniel A Surinach, Ibrahim Oladepo, Ridwan Hossain, Skylar Fausner, Kapil Saxena, Eunsong Ko, Ryan Peters, Michael Feldkamp, et al. Pan-cortical cellular imaging in freely behaving mice using a miniaturized micro-camera array microscope (mini-mcam). *Science Advances*, 11(44):eadt3634, 2025.

[6] Arun Cherkkil, Zoey Viavattine, Vamsy Kota, Kapil Saxena, Daniel Surinach, Ibrahim Oladepo, Jill Juneau, Jia Hu, Skylar M.L. Fausner, Eunsong Ko, Shubham Panchal, Sanjana Srivatsa, Malachi R. Lehman, Roarke Horstmeyer, and Suhasa B Kodandaramaiah. Cortex-wide cellular imaging in freely locomoting mice using cortex camera array microscope (cortexcam). *bioRxiv*, 2026.

[7] Yuanlong Zhang, Lekang Yuan, Qiyu Zhu, Jiamin Wu, Tobias N"obauer, Rujin Zhang, Guihua Xiao, Mingrui Wang, Hao Xie, Zengcai Guo, Qionghai Dai, and Alipasha Vaziri. A miniaturized mesoscope for the large-scale single-neuron-resolved imaging of neuronal activity in freely behaving mice. *Nature Biomedical Engineering*, 8(6):754–774, 2024.

[8] Changliang Guo, Garrett J Blair, Megha Sehgal, Federico N Sangiuliano Jimka, Arash Bellafard, Alcino J Silva, Peyman Golshani, Michele A Basso, Hugh Tad Blair, and Daniel Aharoni. Miniscope-lfov: A large-field-of-view, single-cell-resolution, miniature microscope for wired and wire-free imaging of neural dynamics in freely behaving animals. *Science advances*, 9(16):eadg3918, 2023.

[9] Pingping Zhao, Changliang Guo, Mian Xie, Liangyi Chen, Peyman Golshani, and Daniel Aharoni. MiniXL: An open-source, large field-of-view epifluorescence miniscope enabling single-cell resolution and multi-region imaging in mice. *Science Advances*, 11(24):ads4995, 2025.

[10] Jesse K Adams, Dong Yan, Jimin Wu, Vivek Boominathan, Sibo Gao, Alex V Rodriguez, Soonyoung Kim, Jennifer Carns, Rebecca Richards-Kortum, Caleb Kemere, et al. In vivo lensless microscopy via a phase mask generating diffraction patterns with high-contrast contours. *Nature Biomedical Engineering*, 6(5):617–628, 2022.

[11] Jimin Wu, Yuzhi Chen, Ashok Veeraraghavan, Eyal Seidemann, and Jacob T Robinson. Mesoscopic calcium imaging in a head-unrestrained male non-human primate using a lensless microscope. *Nature Communications*, 15(1):1271, 2024.

[12] Josefa R. Scherrer, Galen F. Lynch, Jie J. Zhang, and Michale S. Fee. An optical design enabling lightweight and large field-of-view head-mounted microscopes. *Nature Methods*, 20(4):546–549, 2023.

[13] Feng Tian, Ben Mattison, and Weijian Yang. Deepinminiscope: Deep learning–powered physics-informed integrated miniscope. *Science Advances*, 11(37):eadr6687, 2025.

[14] Yujia Xue, Ian G Davison, David A Boas, and Lei Tian. Single-shot 3d wide-field fluorescence imaging with a computational miniature mesoscope. *Science advances*, 6(43):eabb7508, 2020.

[15] Minho Eom, Seungjae Han, Pojeong Park, Gyuri Kim, Eun-Seo Cho, Jueun Sim, Kang-Han Lee, Seonghoon Kim, He Tian, Urs L Böhm, et al. Statistically unbiased prediction enables accurate denoising of voltage imaging data. *Nature Methods*, 20(10):1581–1592, 2023.

[16] Pengcheng Zhou, Shanna L Resendez, Jose Rodriguez-Romaguera, Jessica C Jimenez, Shay Q Neufeld, Andrea Giovannucci, Johannes Friedrich, Eftychios A Pnevmatikakis, Garret D Stuber, Rene Hen, et al. Efficient and accurate extraction of in vivo calcium signals from microendoscopic video data. *elife*, 7:e28728, 2018.

[17] Andrea Giovannucci, Johannes Friedrich, Pat Gunn, Jérémie Kalfon, Brandon L Brown, Sue Ann Koay, Jiannis Taxidis, Farzaneh Najafi, Jeffrey L Gauthier, Pengcheng Zhou, et al. Caiman an open source tool for scalable calcium imaging data analysis. *elife*, 8:e38173, 2019.