bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

A deep learning approach for improving two-photon vascular imaging speeds

Annie Zhou[1], Samuel A. Mihelic[1], Shaun A. Engelmann[1], Alankrit Tomar[1], Andrew K. Dunn[1+] & Vagheesh M. Narasimhan[2,3+]

[1]*Department of Biomedical Engineering, The University of Texas at Austin, 107 W. Dean Keeton C0800, Austin, TX 78712, USA*

[2]*Department of Integrative Biology, The University of Texas at Austin, 2415 Speedway C0930, Austin, TX 78712, USA*

[3]*Department of Statistics and Data Sciences, The University of Texas at Austin, 105 E. 24th St D9800, Austin, TX 78712, USA*

[+]Co-corresponding authors

Abstract

A potential method for tracking neurovascular disease progression over time in preclinical models is multiphoton fluorescence microscopy (MPM), which can image cerebral vasculature with capillary-level resolution. However, obtaining high-quality, three-dimensional images with a traditional point scanning MPM is time-consuming and limits sample sizes for chronic studies. Here, we present a convolutional neural network-based algorithm for fast upscaling of low-resolution or sparsely sampled images and combine it with a segmentation-less vectorization process for 3D reconstruction and statistical analysis of vascular network structure. In doing so, we also demonstrate that the use of semi-synthetic training data can replace the expensive and arduous process of acquiring low- and high-resolution training pairs without compromising vectorization outcomes, and thus open the possibility of utilizing such approaches for other MPM tasks where collecting training data is challenging. We applied our approach to large field of view images and show that our method generalizes across imaging depths, disease states and other differences in neurovasculature. Our pre-trained models and lightweight architecture can be used to reduce MPM imaging time by up to fourfold without any changes in underlying hardware, thereby enabling deployability across a range of settings.

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Introduction

The neurovascular network transports chemicals (e.g., oxygen, nutrients, waste) to and from the brain to support neuronal activity. Neurovascular function is disrupted by disorders such as stroke, Alzheimer's and other neurodegenerative diseases, and diabetes, with lasting effects that are not fully understood[1,2]. Advances in multiphoton fluorescence microscopy (MPM) have enabled imaging with capillary-level resolution *in vivo*, and this noninvasive tool could be used to monitor capillary-level changes over time in cerebral vasculature as a potential predictor of disease progression/prognosis[3–6]. A constraint with MPM, however, is the slow acquisition process that is necessary for producing high quality, three-dimensional images with a traditional point scanning multiphoton imaging setup. Given the physical limitations of a live animal, the acquisition speed puts a limit on study sizes and the ability to reach statistically significant results. Although previous approaches have sought to improve imaging speeds by incorporating innovative imaging hardware, these implementations come at high cost and complexity and cannot be readily employed to existing infrastructure[7–10].

An alternative, more cost-effective and accessible approach might be to computationally improve the image acquisition process using Convolutional Neural Networks (CNNs), which leverage existing datasets of MPM images. While several recent advances have been made in applying CNNs to improve MPM imaging results[11–16], to our knowledge, only one has been focused on improving MPM imaging speed. Guan et al. presented a CNN for improving the imaging speed of a two-photon fiberscope for neuronal imaging using a conditional generative adversarial network (cGAN)[16]. They achieved a 10-fold speedup in frame rate. A drawback to their approach, however, is the requirement for a two-part training set, involving both *ex vivo* and *in vivo* imaging, which is extremely expensive and time-consuming. Several other models for

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

general denoising or segmentation for MPM have also been focused primarily on neuronal or calcium imaging[11–14], with only one to our knowledge focused on vascular segmentation[15]–none of which is used for improving vascular imaging speeds.

Here, we describe a CNN based approach trained to take images captured at low-resolution (128 x 128 pixels), thereby at much faster speeds, and then upscale these to a higher resolution (512 x 512 pixels), without compromising the accuracy of vascular morphological information that is extracted or introducing additional noise. The upscaling process from low- to high- resolution using deep learning is referred to as image super-resolution. We then combined this with a vectorization pipeline to obtain quantitative statistics of neurovasculature. Our pre-trained models and light architecture allows for fast acquisition, image super-resolution, and vectorization of MPM images without the limitations of added hardware and can be used to reduce imaging time by up to fourfold.

Results

Structure and analysis pipeline overview

Our pipeline to improve two-photon microscopy acquisition and vectorization accepts individual as well as multiple frames from MPM imaging. Our process is split into two main parts: an image super-resolution CNN designed to upscale low-resolution images and a vectorization pipeline that is designed to output morphological statistics on the acquired images (**Fig. 1**). For super-resolving the images, we used the PSSR Res-U-Net architecture, which had been shown to restore images of presynaptic vesicles and neuronal mitochondria from a scanning electron microscope (SEM) and a laser scanning confocal microscope, respectively[17]. We utilized this architecture over several other possible ones because it: (a) allowed us to use semi-synthetic data

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

for training which circumvents the need to acquire real-world image pairs for training, which is difficult and expensive for large datasets (b) enabled us to also employ multi-frame inputs that could leverage information across correlated images at similar depths (c) does not utilize an adversarial network in the training, which are more challenging to train as well as to evaluate the generated models (d) allowed us to use a transfer learning approach to initialize our model with weights obtained with the architecture trained on ImageNet, a large natural image classification dataset. For vectorization, we used Segmentation-Less, Automated, Vascular Vectorization (SLAVV)[18], which uses simple models of vascular anatomy, and efficient linear filtering and vector extraction algorithms with manual or automated vector classification. Using a multi-frame PSSR approach and combining it with a vectorization pipeline, we show that we are able to restore two-photon vascular images sufficiently for extracting morphological characteristics through vectorization.

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

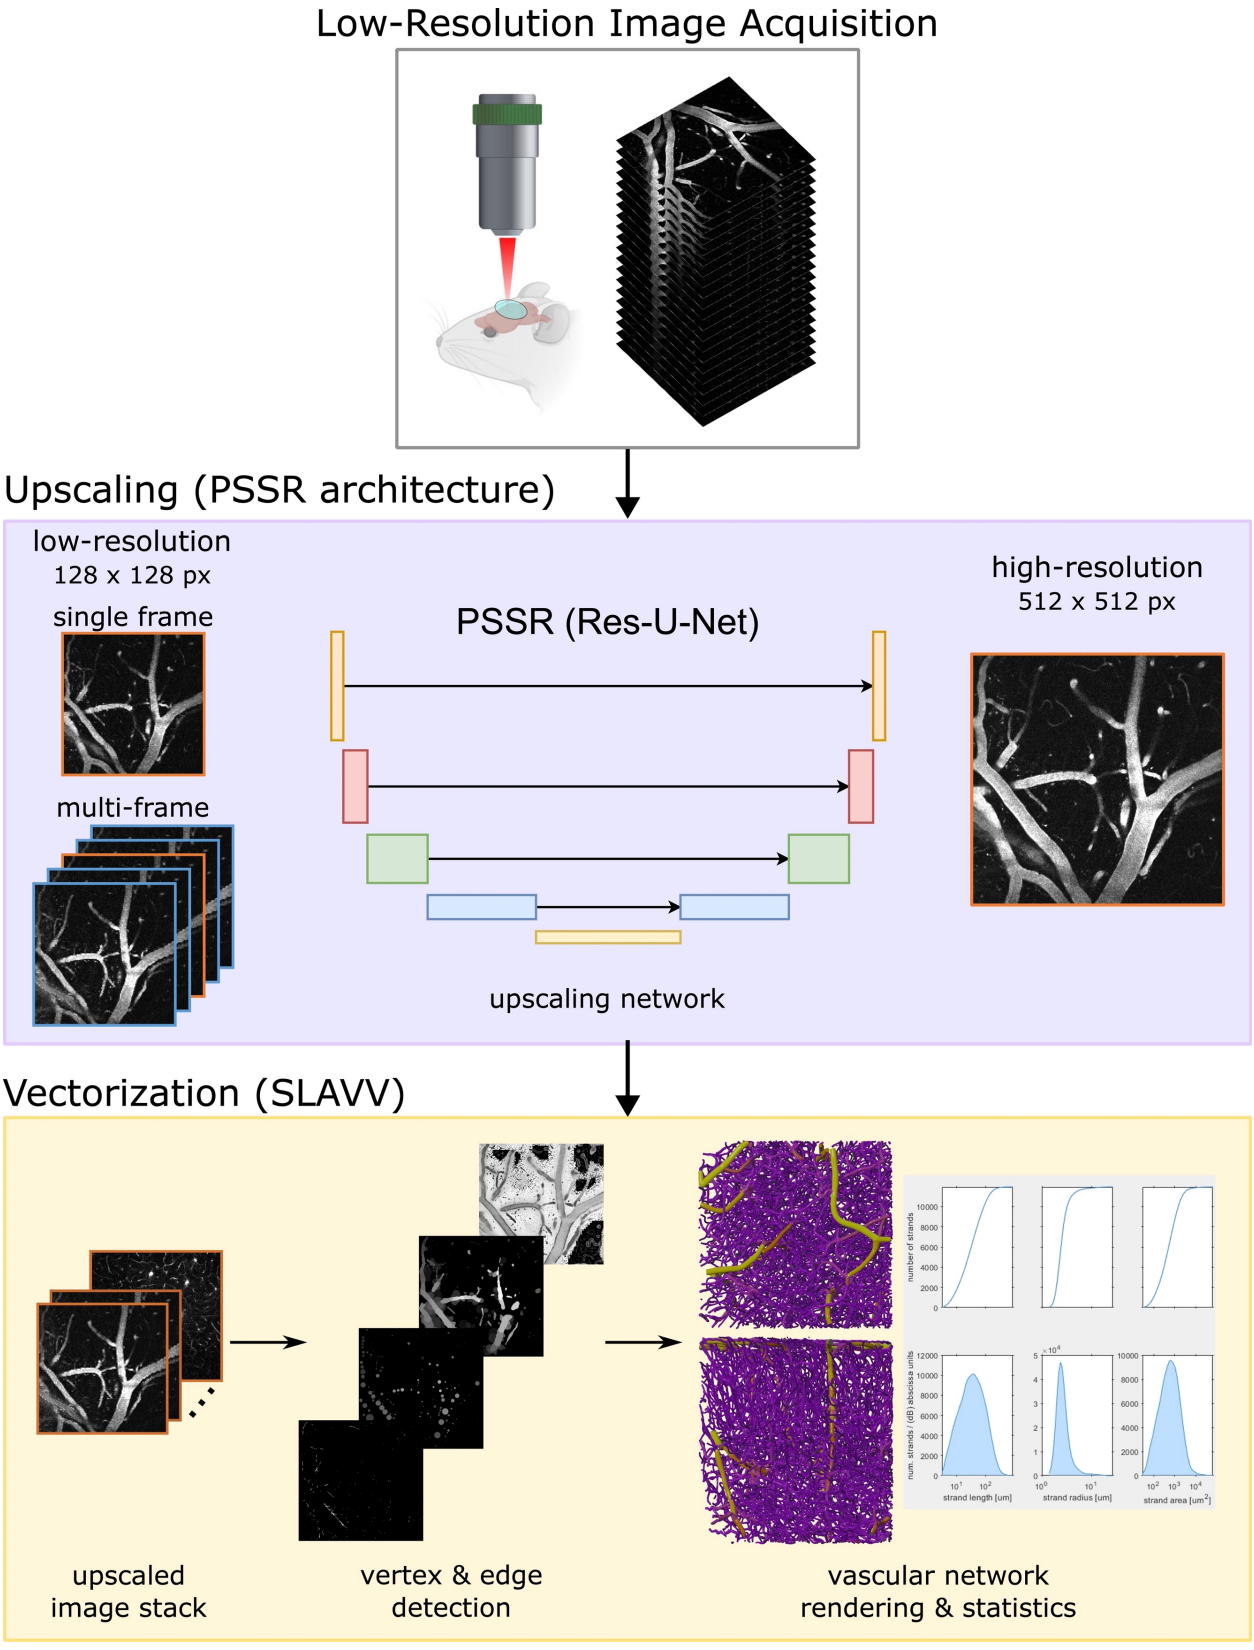

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Fig. 1: Structure and analysis pipeline. Low-resolution images (128 x 128 pixels) are acquired using two-photon microscopy. A deep learning (PSSR Res-U-Net) based upscaling process generates high-resolution images (512 x 512 pixels), which take much longer to acquire, from low-resolution images. Segmentation-less vascular vectorization (SLAVV) generates 3D renderings and calculates network statistics from an upscaled image stack.

Transfer learning, creation and evaluation of semi-synthetic training data

Traditional approaches to upscaling images involve acquiring paired high- and low-resolution real-world images that we could use for training the model. For our task, however, this process is expensive and time-consuming, and in certain situations impossible for live animals. This challenge greatly limits the practical sample size of training datasets. To overcome this difficulty, we sought to use semi-synthetic training data that mimics low-resolution acquisition to greatly improve sample size. Semi-synthetic training data was created by adding noise to, then downscaling, full-resolution images from a two-photon vascular image repository of previously acquired images (see Data Availability). We evaluated several approaches for the creation of this semi-synthetic data and compared our results to a model trained with a real-world dataset of the same sample size.

To mimic the noise observed in real-acquired low-resolution images (i.e., real data), we tested models trained with semi-synthetic images that were created with the following noise filters: no noise (downscaled only, used as the reference), Poisson noise, Gaussian noise, and additive-Gaussian-distributed noise (**Fig. 2a**). Real-acquired low-resolution images served as input images to test the model. We evaluated model performance using standard image quality metrics,

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

specifically, by calculating peak signal-to-noise ratio (PSNR) and structural similarity index measure (SSIM) between the model output and acquired full resolution image.

The resulting median PSNR and SSIM values from each model, ranked from highest to lowest for both metrics, were as follows: Gaussian, additive-Gaussian, Poisson, no noise (**Fig. 2b-c**). This was determined using Wilcoxon signed rank tests with $P < 0.005$ (Bonferroni-adjusted). We noticed that the Gaussian and additive-Gaussian models perform similarly, and thus performed further testing to compare the two noise methods using a larger training set, consisting of 24,069 semi-synthetic training image pairs—7x the preliminary training set of 3,399 semi-synthetic image pairs. The test image outputs from the models trained with the larger dataset showed notable qualitative improvements, with fewer false detections, less noise, and smoother vessel shapes. With the larger dataset, the Gaussian model produced a slightly higher median PSNR value but did not produce a median SSIM value that was statistically significantly different from that produced by additive-Gaussian (Wilcoxon signed rank test, $P < 0.005$). Despite the slightly higher PSNR performance by the Gaussian model, however, a qualitative comparison suggested that the additive-Gaussian results had somewhat less noise and higher sensitivity to fainter vessels. Additionally, PSNR only measures similarity in pixel values and does not necessarily predict vectorization performance, which is what we ultimately wish to optimize. To fully validate and compare the performance between the Gaussian- and additive-Gaussian-trained models, we performed a final comparison test using the Segmentation-Less Automated Vascular Vectorization (SLAVV) software (further described in later section). We found that the additive-Gaussian model output allowed for a more accurate vessel detection overall compared to the Gaussian model output (95.7% vs 95.6%). Based on these results, we chose to perform

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

subsequent analyses using the model trained with the full dataset of semi-synthetic images created using the additive-Gaussian noise method to maximize accurate vessel detection.

To evaluate the effectiveness of using semi-synthetic data in place of real-world training data, we compared the performance of models trained with each method (**Fig. 2d-e**). For this comparison, both models were trained with 234 image pairs, due to the limited availability of acquired image pairs. The output image from the real-acquired model appeared blurrier and over-predicted vessel diameters more significantly compared to the semi-synthetic model (**Supplemental Figure 1**). Nonetheless, the model trained with real-world data had higher median PSNR and SSIM values compared to the model trained with semi-synthetic data (Wilcoxon signed rank test, $P < 0.05$), although the values were close (PSNR: 26.9 vs. 26.6; SSIM: 0.492 vs. 0.494). We deem the results similar enough for semi-synthetic training data to be used in place of real-world training data. The use of semi-synthetic data advantageously circumvents complications from imprecise alignment in the acquisition of image pairs, limited availability of existing images (677 pairs), and high material and labor costs for data collection.

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

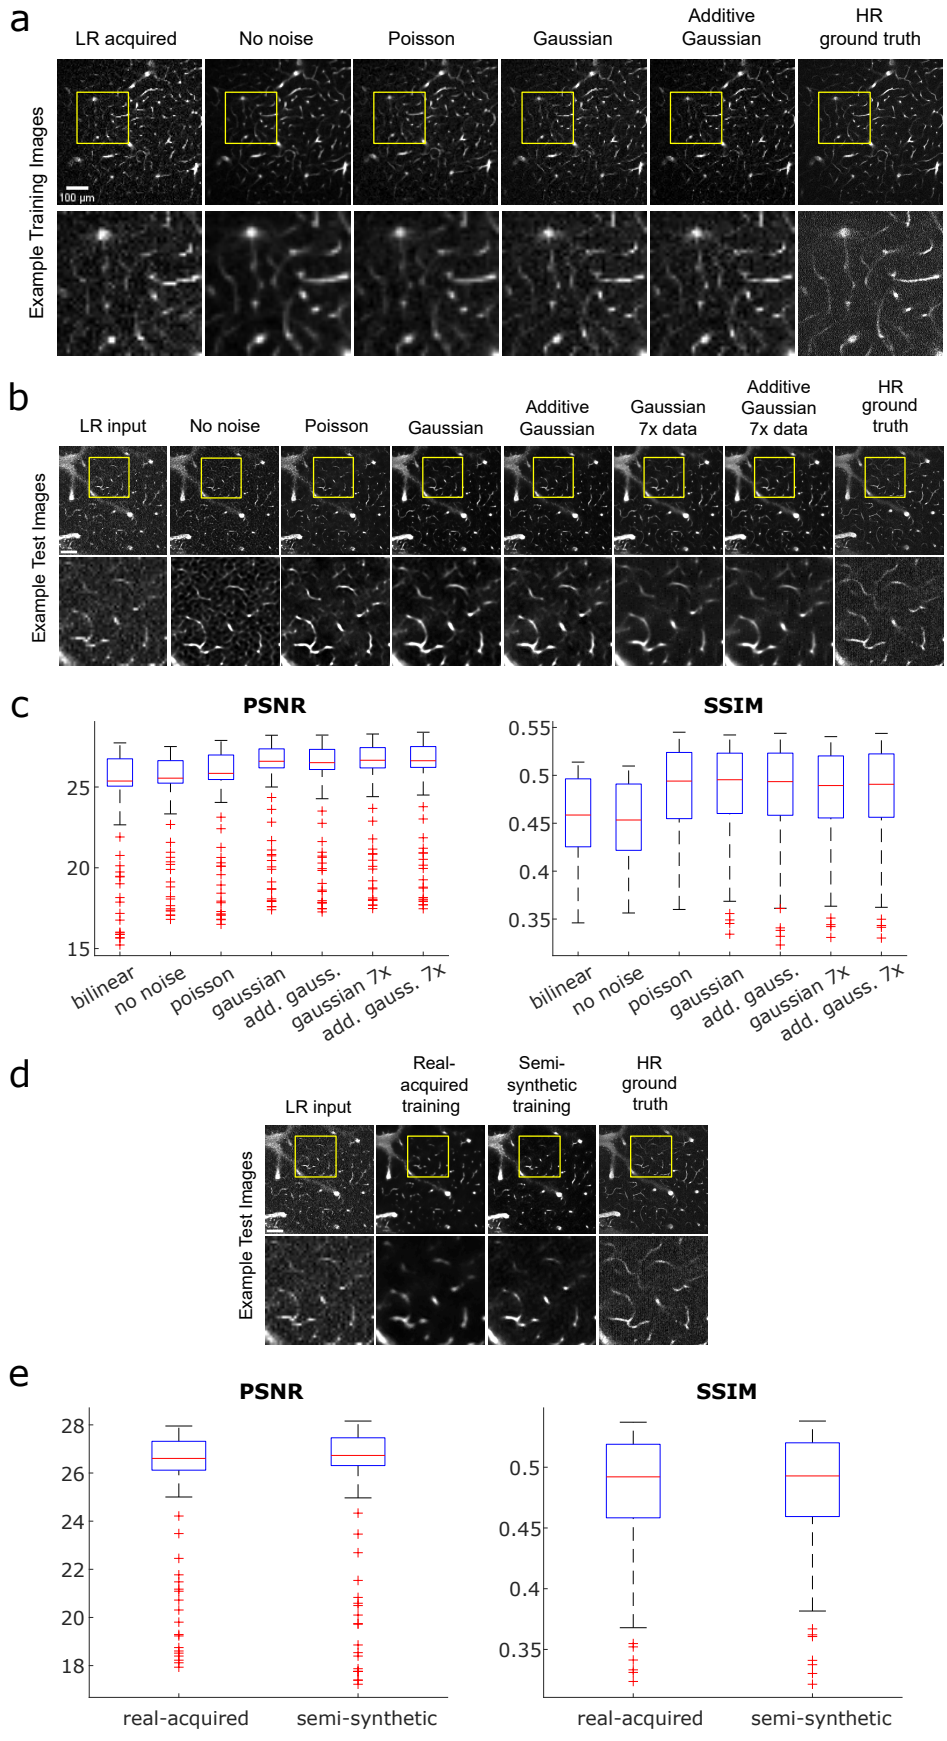

139

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Fig. 2: **Generating and evaluating semi-synthetic training data. a** Examples of semi-synthetic training images created using different types of added noise prior to downscaling: no noise (downscaling only), Poisson, Gaussian, and additive Gaussian. Acquired low-resolution (LR, 128x128 pixels) and high-resolution (HR, 512x512 pixels) ground truth images are shown for reference. **b** Resulting test images from models trained using each noise method, with acquired low-resolution image for model input and acquired high-resolution image as ground truth for comparison. All models were trained with 3,399 image pairs, with the Gaussian and additive Gaussian models further tested on 24,069 image pairs (7x) to further test performance. **c** Boxplot comparison of PSNR and SSIM values for each noise method image in **b** measured against ground truth image. Values plotted for an image stack of 222 images. **d** Comparison of test images from models trained using real-world acquired vs. semi-synthetic data, with real-acquired low-resolution image for model input and acquired high-resolution image as ground truth for reference. All models were trained with 234 image pairs, a large reduction from the noise model comparison, due to the limited availability of real-world pairs. **e** Boxplot comparison of PSNR and SSIM values for real-acquired vs. semi-synthetic model outputs corresponding to **d**, measured against high-resolution ground truth image. All values are plotted for an image stack of 222 images.

Single-frame vs. multi-frame training

A key issue with low-resolution acquired images is the diminished amount of total signal capture. This can result in noisier images and cause spurious vessels to appear in the vectorization process. A potential method for reducing false detections is providing the model neighboring depth images on a stack, which are highly correlated in signal but not in noise. Thus, we sought to improve the performance of our model by using multi-frame image input.

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

We compare the performance of the single-frame model to a multi-frame model, with an additional comparison against the traditional bilinear upscaling method, for both semi-synthetic and real-world test images (**Fig. 3**). The traditional bilinear upscaling method offers a baseline performance measure for a non-CNN approach. The multi-frame model is trained with input image stacks consisting of five sequential images in depth, with axial offsets of 0.3 μm, to predict a single output image—the third image in the input sequence. The multi-frame model yields images with higher PSNR and SSIM values than the single-frame model, and both PSSR methods outperform the bilinear upscaling method for both semi-synthetic and real-world test images (Wilcoxon signed rank test, $P < 0.0167$, Bonferroni-adjusted). Overall PSNR and SSIM values are higher for semi-synthetic images compared to real-world images, which is unsurprising given the model was trained completely on semi-synthetic images. Nonetheless, the real-world output images from our models show that individual vessels can be resolved, which is much more important for the final vectorization process than the exact pixel values measured by PSNR.

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

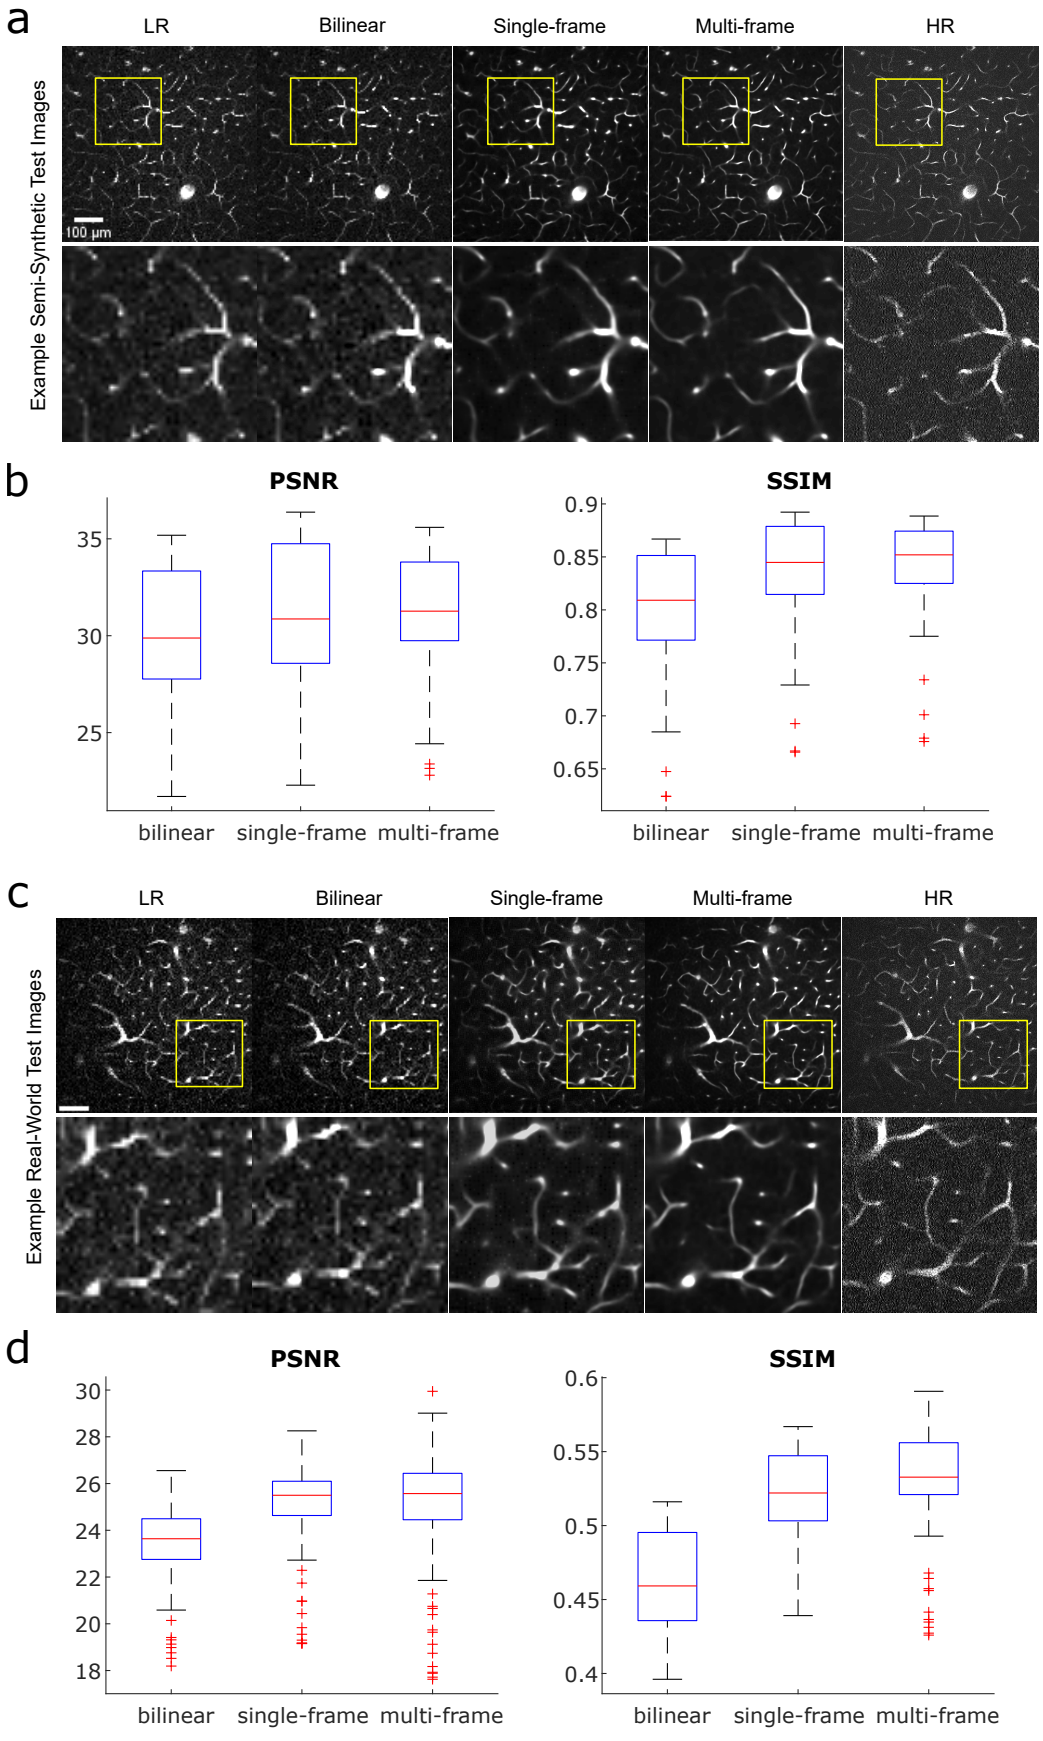

177

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Fig. 3: Comparison of performance between bilinear upscaling, a single-frame model, and a multi-frame model for semi-synthetic and real-acquired test images. All models were trained with 24,069 image pairs. **a** Semi-synthetic test images from bilinear upscaling and models trained using single- vs. multi-frame data. Acquired low-resolution image for model input and acquired high-resolution image as ground truth are shown for reference. **b** PSNR and SSIM plots corresponding to semi-synthetic test results from **a**. **c** Real-world test images from bilinear upscaling and models trained using single- vs. multi-frame data. **d** PSNR and SSIM plots corresponding to real-world test image results from **c**.

Reconstruction and stitching of infarct images

A major application of two-photon imaging that we aim to make more accessible with our approach is large field of view (FOV) imaging of diseased vasculature. An example of this is acquiring images from a stroke model, which is of interest for studying disease effects on vascular morphology. Large FOV imaging with high resolution is a time-consuming process and thus would benefit substantially from the speedup offered by low-resolution imaging. Large FOV images are achieved by acquiring, then stitching standard-sized tiles together using ImageJ's Grid/Collection Stitching plugin[19]. However, the long acquisition times required of current technologies limit our ability to collect substantial two-photon image sets of diseased vascular networks and thus limit the availability of images of diseased vasculature that could be used for training data.

To investigate the feasibility of using our models to drastically reduce imaging times for large FOV images of diseased vasculature with minimal information loss, we examined the ability of our single-frame and multi-frame models, trained only with images of normal vasculature, to

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

restore a semi-synthetic large FOV image of an ischemic infarct (four weeks post-stroke) collected in a preliminary study (**Fig. 4**). The differences in morphology between vasculature in the peri-infarct region and normal vasculature are exemplified by the differences between the top half of the full image, which more closely resembles normal vasculature, and the bottom half of the image, which captures the infarct region and the more immediately surrounding vessels. Ischemic infarct vessels appear significantly more parallel to the imaging plane, thus creating image slices with higher vascular area density compared to the more perpendicularly oriented vessels further from the infarct. Despite these morphological differences and having only trained with images of normal vasculature, our models are able to resolve capillaries in the infarct region, as shown in the insets of **Fig. 4**. The multi-frame output image more closely resembles the HR image compared to the single-frame image, as vessel radii are more consistent in the multi-frame image. In the case of the LR and bilinear-upscaled images, the individual capillaries in the inset cannot be resolved.

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

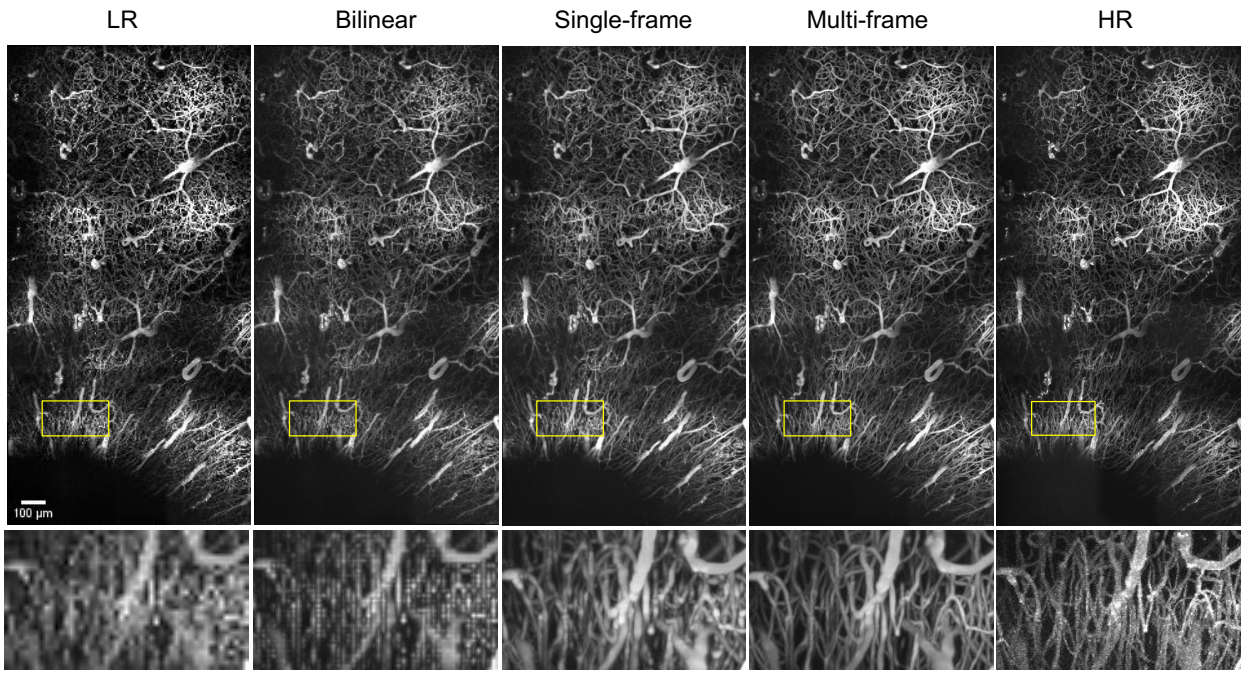

Fig. 4. Maximum intensity projections (x-y) of ischemic infarct images consisting of 2x4 tiles with 213 slices (final dimensions 1.18 mm x 2.10 mm x 0.636 mm, pixel dimensions 1.34 µm x 1.36 µm x 3 µm) for a semi-synthetic low-resolution image, bilinear upscaled image, single- and multi-frame output images, and acquired high-resolution image.

Vectorization

Vectorization is the ultimate step that extracts quantitative information for evaluating the vascular morphology of a network. Therefore, we are interested in comparing different image generation strategies by comparing performance after vectorization. We demonstrate successful vectorization of single- and multi-frame model output images from real-acquired low-resolution images using manual-curation-assisted SLAVV and visualization with VessMorphoVis[20] (**Fig. 5a**). Additionally, we perform a more objective comparison of our models' performance using a previously described method[18], which uses simulated images from a known ground truth and

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

automated vector classification (no manual assist). Using a known ground truth (derived from the real-acquired high resolution vectorized network shown in (**Fig. 5a**), we are able to quantify the sensitivity, specificity, and accuracy of the vectorized upscaled images. We generated the simulated vascular image to have the same contrast-to-noise ratio (CNR) of 0.94 as the real acquired high resolution image and to be representative of the image quality of a typical image acquired by our two-photon microscope. We created a low-resolution version of the simulated image using the same method for creating semi-synthetic training data and then upscaled it using bilinear interpolation and our single-frame and multi-frame models. We then vectorized these images using fully automated (globally thresholded) SLAVV at peak segmentation performance (measured against the ground truth image). The resulting strand objects are the minimal set of one-dimensional traces which span the entire vascular network.

We plotted cumulative distribution functions (CDFs) for each image for each of the following strand metrics: length, radius, z-direction, and inverse tortuosity (**Fig. 5b**). We included the simulated original image in the analysis as a control for the automated curation process since the ground truth image was obtained through manual curation. For each strand metric, we calculated Pearson's correlation (r^2) values between the CDFs of the ground truth image and the simulated images. Of all the images, the simulated original image maintained the highest r^2 value for average strand radius and inverse tortuosity. Our multi-frame model had the highest r^2 value for strand length, while bilinear and the single-frame model produced the highest r^2 value for z-direction. We performed a Kolmogorov–Smirnov (K-S) test to compare the CDFs of each upscaling method against that of the simulated original image. The multi-frame CDFs for strand length, radius, and z-direction; the single-frame CDFs for length and z-direction; and the bilinear CDF for z-direction were not significantly different from those of the simulated original image (p

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

250 < 0.0167, Bonferroni-adjusted). Thus, of the tested upscaling methods, the multi-frame model
251 produced the most statistically comparable strand metrics to the simulated original image.

252 We calculated overall accuracy with respect to the ground truth for each image (**Fig. 5c**). The
253 original simulated image retains the highest vectorization accuracy (96.2%), followed by multi-
254 frame (95.7%), single-frame (95.2%), and bilinear (94.5%). In terms of accuracy with raw image
255 segmentation through intensity thresholding, however, multi-frame performs best (96.0%),
256 followed by single-frame (95.3%), bilinear (94.4%), and original simulated image (91.9%). We
257 also calculated the percent error in the median length and median radius, the characteristics that
258 best represent the vessel morphology, between each image against the ground truth values.
259 Multi-frame produced the lowest median length error (6.4%), followed by bilinear (7.2%),
260 single-frame (8.1%), and the original simulated image (8.2%). The bilinear and single-frame
261 images had notably higher median radius errors (40.1% and 41.6%, respectively) compared to
262 the multi-frame and original simulated images (both 26.9%), which was noted with visual
263 inspection of the images as well. These statistics further support that a multi-frame upscaled
264 image produces comparable vectorization results to an original high-resolution image.

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

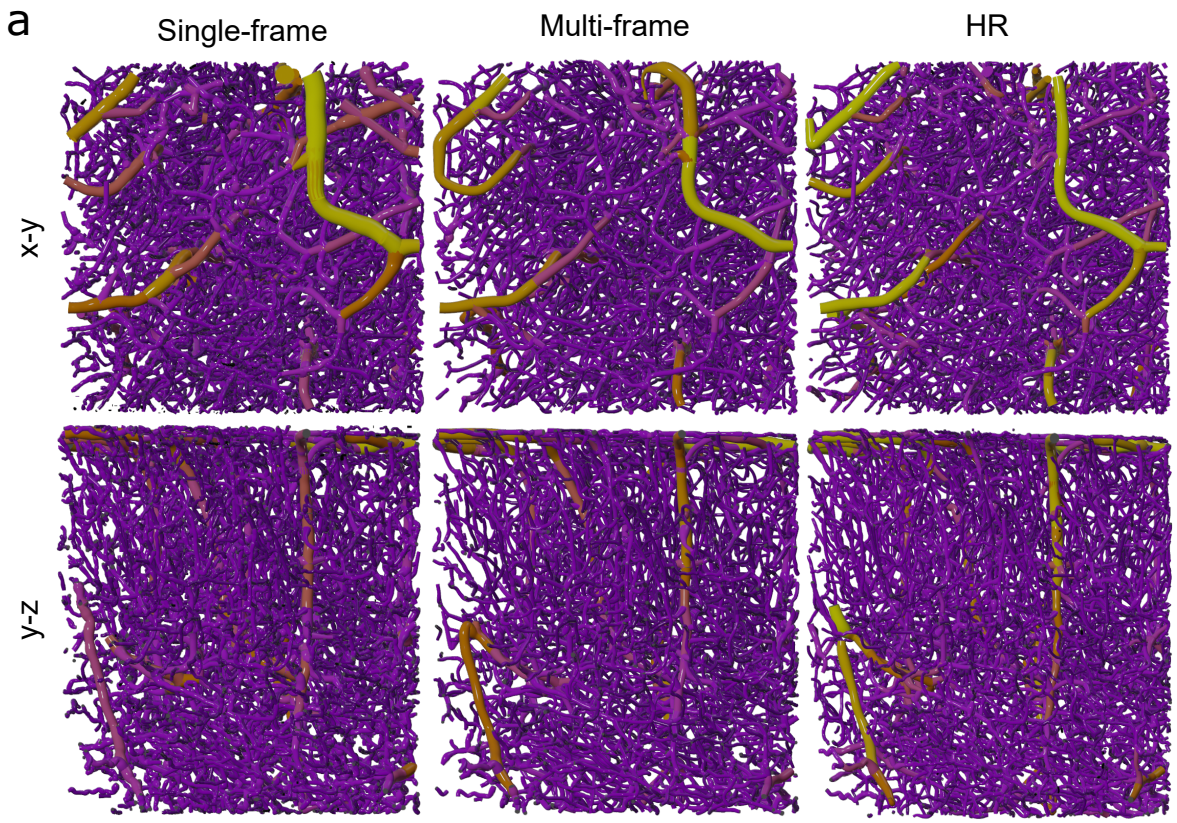

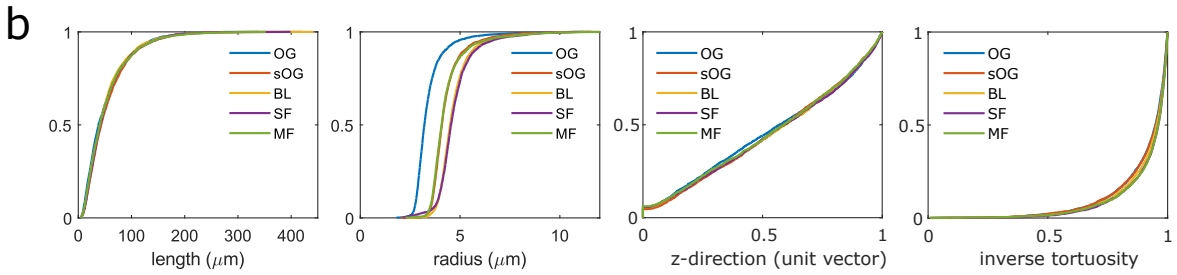

r^2	length	radius	z-direction	inv. tortuosity
Simulated original (sOG)	0.9994	0.9714	0.9991	0.9999
Bilinear upsample (BL)	0.9981	0.9049	0.9994	0.9993
PSSR Single-Frame (SF)	0.9993	0.9054	0.9994	0.9973
PSSR Multi-Frame (MF)	0.9997	0.9654	0.9992	0.9972

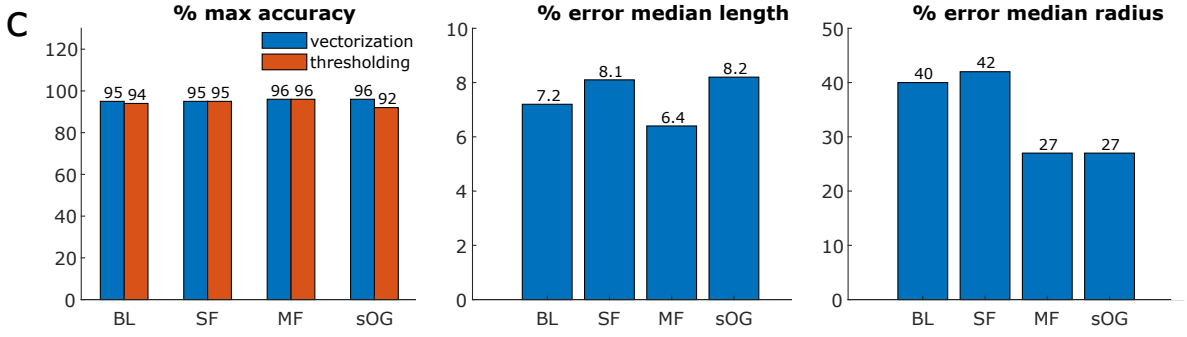

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Fig. 5. Vectorization results. a Blender rendering of vectorized images using VessMorphoVis[20] for visual comparison between single- and multi-frame results and an acquired high-resolution image. We performed manual curation for this vectorization process. **b** Vectorized image statistics for automated curation process with known ground truth (simulated from manually curated high-resolution image). CDFs shown for metrics of length, radius, z-direction, and inverse tortuosity for original (OG), simulated original (sOG), bilinear upscaled (BL), and PSSR single- and multi-frame (SF, MF, respectively) images. Pearson's correlation values (r^2) were calculated between the original image and each simulated or upscaled image for each metric. **c** Statistics regarding maximum accuracy (%) achieved with vectorization or thresholding and % error in median length and radius for each method.

Discussion

To our knowledge, this is the first time that a deep learning model has been demonstrated to improve imaging speeds for two-photon microscopy by upscaling and denoising low-resolution images of vasculature while retaining accuracy in extracted morphological characteristics. For this application, our model outperforms the traditional, non-CNN bilinear upscaling technique in output image quality (**Fig. 3, 4**) and vectorization accuracy (**Fig. 5**). The performance of our models also improves notably with increased training data (**Fig. 2b**), therefore substantial time and material costs are reduced by training the model with semi-synthetic images generated from our database of 28,563 previously acquired two-photon vascular images. Real-world data also introduces further complications by requiring image registration. Not only would this add computational hours, but the image registration process also does not produce perfect alignment because it is limited to being purely translational and free of interpolation to maintain the original recorded pixel values. Any rotational misalignments would not be accurately

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

289 correctable. We speculate that these misalignments in real-world training image pairs caused the
290 overly blurry and enlarged vascular structures seen in the results of preliminary experiments
291 (**Fig. 2d**, **Supplementary Fig. 1**).

292 Models trained from semi-synthetic images proved capable of restoring low-resolution vascular
293 images and outperformed the standard bilinear interpolation method. For performing
294 segmentation via intensity thresholding (**Fig. 5c**), multi-frame had the highest accuracy of the
295 upscaling methods and significantly outperformed the original standard-resolution image. We
296 hypothesize that this is a result of the denoising that occurs in the PSSR process. Since all
297 upscaled images had higher intensity thresholding accuracy compared to the original image, we
298 further postulate that upscaled images have less noise overall because fewer pixels are physically
299 captured—all pixels created during upscaling are interpolated from neighboring pixel values and
300 thus free of noise from the image acquisition process[21]. For performing vectorization, however,
301 the diminished noise does not offset accuracy losses from the upscaling process. We determined
302 that the multi-frame model yields the highest accuracy of the three upscaling methods but did not
303 outperform the original standard-resolution image. Nonetheless, the multi-frame image also
304 produced the greatest number of CDFs for strand metrics that were not significantly different
305 from those of a standard-resolution image. We consider these results from the multi-frame model
306 to be within acceptable tolerance for vectorization accuracy and similarity in strand statistics for
307 our previously described purposes of characterizing the structural properties of the vessels of a
308 particular network[22,23].

309 By acquiring low-resolution images, the imaging time could be reduced by up to half in a two-
310 photon microscope with a resonant-galvo scanning system and by up to fourfold with a galvo-
311 galvo scanning system. The reduction in imaging time can have several benefits. For instance,

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

faster imaging times reduce risk of phototoxicity and thermal damage if excitation powers are kept constant[24,25]. Additionally, the amount of time for which the subject is under anesthesia is reduced, decreasing the risk of vascular dilation[26] which can create skewed vectorization statistics. The reduction in imaging time can also decrease the injection volume and frequency of fluorescent dye, which alone can save up to hundreds of dollars in addition to eliminate the risk of sample misalignment caused by a re-injection during an imaging session. A specific but major benefit for those wishing to conduct chronic studies is that the faster acquisition times will allow for larger cohorts, which are currently constrained by the number of animals that can reasonably be imaged within each timepoint. This would yield more statistically significant sample sizes for studying and comparing healthy and diseased vasculatures over time.

With the potential for future disease studies in mind, we tested our model on data from a mouse that was given a stroke. We show that our model, despite having only been trained with images of healthy vasculature, can reasonably restore images taken from a peri-infarct region with sufficient resolution for the image stitching algorithm to successfully create a large FOV image. These results from semi-synthetic test data are promising for being able to apply the model broadly to different disease models, although further validation should be performed with real-world test data.

Another area of research that could benefit from increased imaging speeds is in the study of light propagation through the brain, for the development of noninvasive brain imaging devices. With accessibility to a larger database of large FOV two-photon images, light propagation models can be more thoroughly developed, tested, and refined[27]. As precise capillary capture is not necessary for these models, the use of even lower resolution and faster imaging could be further explored with PSSR. Although losses in accuracy in low-resolution images are inevitable due to less

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

information being captured, experimentation with the accuracy and speedup tradeoff can be done to fit the tolerance of any application.

The potential for further speedup by reducing frame averaging could also be explored. Lower frame averaging leads to higher noise levels, which PSSR can be used to reduce. With higher noise, we would expect increased false positive and/or false negative detections, leading to an overall reduction in restoration accuracy. A potential method that can be explored to combat this effect would be to modify the loss function to increase the penalty for false negative detections with the tradeoff of accepting more noise in the image. Alternatively, to prioritize denoising over having high sensitivity, the loss function could be modified to penalize false positive detections more heavily.

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Methods

In vivo imaging

Animal Preparation

Cranial window implants were prepared in C57 mice with dura intact, as previously described[28]. During imaging, mice were anesthetized with isoflurane and body temperature was maintained at 37.5 °C. Blood plasma was fluorescently labeled with dextran conjugated Texas Red (70kDa, D1830, Thermo Fisher) dissolved in saline (5% w/v). The dye was administered intravenously via retro-orbital injection (0.1 mL). All animal protocols were approved by The University of Texas at Austin Institutional Animal Care and Use Committee.

For stroke model mice in **Fig. 4**, photothrombotic ischemia was induced through retro-orbitally injecting rose bengal (0.15ml at 15 mg/ml) and irradiating a penetrating arteriole branching from the middle cerebral artery for 15 minutes. The laser source had a 532 nm wavelength, a 20mW average power, and was focused to a ~300 μm diameter spot size. Mice were anesthetized with isoflurane (1.5%, 0.6-0.8 LPM) and body temperature was maintained with a heating pad during the procedure. Pial anatomy was visualized using laser speckle contrast imaging to select which artery to target and to confirm occlusion.

Image Acquisition

All images were acquired using a custom-built two-photon microscope, previously described[21]. The excitation source was an ytterbium fiber amplifier with an output beam of 1050 nm wavelength, 120 fs pulse width, and 80 MHz repetition rate[29]. High-resolution images were 512x512 pixels and low-resolution images were 128x128 pixels, both with a field size of 700x700 μm. Image stacks were acquired with 3 μm axial spacing. A resonant-galvanometer

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

scanning system was used[21], with average pixel dwell times of 87.8 ns and 20-frame averaging. Power at the sample did not exceed 170 mW and was identical between low- and high-resolution pairs. Large field of view images were taken as a 2-by-4 grid of standard images, with ~25-30% overlap between tiles.

Image processing

Image pre-processing

All images were normalized prior to use as a training image or semi-synthetic test image. A 3D median filter of size [1 1 1] was applied to raw image stacks, followed by a full-scale contrast stretch (FSCS) to fill the 16-bit range with 0.3% saturation across the entire stack, using the normalization function provided by Fiji ImageJ[30]. This FSCS normalization method was determined to create the best images compared to FSCS across the entire stack without saturation and FSCS by image slice (**Supplementary Fig. 2**). Images were then converted from 16-bit to 8-bit and separated into individual frames.

Stitching

Large field of view images acquired as a 2-by-4 grid of standard images were stitched together using ImageJ's Grid/Collection Stitching plugin[19].

Semi-synthetic image generation

Single-frame images

To create semi-synthetic low-resolution images for training, pre-processed real-world images received one of the following types of noise: Poisson, Gaussian ($\mu = 0$, $\sigma = 0.1$), additive-Gaussian ($\mu = 0$, $\sigma = 5$), or no noise prior to fourfold downscaling (from 512 x 512 pixels to 128

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

x 128 pixels). For additive-Gaussian noise, the local variance was scaled by 0.001. A range of parameters (i.e., mean, standard deviation, local variance) were tested to identify values for optimal performance. Models were trained for each combination of parameters and given test images. Output images were inspected visually, and image quality metrics (PSNR, SSIM) were calculated.

Multi-frame images

Low-resolution images from the single-frame semi-synthetic image generation step were used to create multi-frame training images. Multi-frame images consisted of five low-resolution images sequential in axial space with 0.3 μm separation (axial distance between acquired images).

Neural networks and training

The Res-U-Net architecture as described by Fang et al. as PSSR (point-scanning super resolution) was used for single-frame and multi-frame training[17]. We used an MSE loss function after determining that L1 and feature loss did not perform as well (**Supplementary Fig. 3**). A learning rate of 9e-4 was used for single-frame training and 1e-4 was used for multi-frame training.

Training/Test Images

Preliminary training for finding the best noise model was done with 3,399 training (data from 5 mice, 16 stacks, 6 imaging sessions) and 676 validation image pairs (2 mice, 3 stacks, 2 imaging sessions). Final full dataset training was completed using 24,069 training (6 mice, 114 stacks, 19 imaging sessions) and 4,494 validation (6 mice, 22 stacks, 7 imaging sessions) image pairs. Real-acquired image pairs (677 image pairs from 2 mice, 3 stacks, 2 imaging sessions) were used for testing and evaluating the models. These image pairs were also used for comparing the

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

411 performance between training with real-acquired pairs vs. semi-synthetic pairs (234 image pairs
412 for training, 221 image pairs for validation, 222 image pairs for testing; each set from 1 mouse, 1
413 stack, 1 imaging session).

414 *Hardware*

415 Training was performed using Frontera at the Texas Advanced Computing Center (TACC) with
416 four NVIDIA Quadro RTX 5000 GPUs using the CUDA version 10.0 toolkit.

417 ***Image Quality Evaluation***

418 Image quality between upscaled images and the original image was preliminarily assessed with
419 peak signal-to-noise ratio (PSNR) and structural similarity index measure (SSIM). Both metrics
420 were computed using built-in MATLAB functions. In combination, these metrics gave a general
421 sense for image similarity, but were not indicators for morphological accuracy from
422 vectorization. Generally, higher PSNR and SSIM values are desired. However, both metrics
423 correlate similarities in raw intensity values to higher similarity between images, whereas the
424 vectorization process is designed for vascular networks, is the end-goal of image acquisition, and
425 produces quantitative anatomical information which may or may not be of interest to a particular
426 researcher. For simplicity, when possible, we chose image segmentation accuracy (with respect
427 to the ground truth) to measure general vectorization performance. In the case of unknown
428 ground truth, SLAVV was used to segment the original image and estimate the CNR, which was
429 used to match the quality of the simulated and real-acquired images. PSNR and SSIM values do
430 not seem to fully reflect image quality improvement from denoising, as seen in **Fig. 2b** where
431 image quality improves visually with the increased training data, but PSNR and SSIM values
432 decrease slightly. In addition, higher variations in predicted pixel intensity value are seen within

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

433 white vessel regions, which can cause lower PSNR and SSIM values despite not affecting visual
434 quality or vectorization performance. This especially affects images closer to the surface of the
435 brain, where large arteries dominate the image, and accounts for the outlier points seen in the
436 boxplots of **Fig. 2** and **Fig. 3**.

437 ***Vectorization***

438 All vectorization was performed using SLAVV software[18]. Real acquired low-resolution images
439 were upscaled using PSSR, then vectorized and manually curated. To obtain an objective
440 comparison of methods without manual curation bias, a simulated image with known ground
441 truth was created from the manually curated high-resolution image. The simulated image had an
442 identical CNR (0.94) to the original acquired high-resolution image, as measured by SLAVV as:

443 $$CNR = \frac{(\mathrm{I}_{foreground} - \mathrm{I}_{background})}{std(\mathrm{I}_{background})}$$

444 where $\mathrm{I}_{\mathrm{foreground}}$ is average foreground intensity and $\mathrm{I}_{\mathrm{background}}$ is average background intensity.

445 A low-resolution image was created using the previously described semi-synthetic image
446 generation method. The low-resolution image was upscaled using bilinear interpolation, the
447 single-frame model, and the multi-frame model. Vectorization of simulated images using
448 automated curation was possible with the known ground truth, as previously described[18].

449 From the vectorized networks for each upscaled image, the original simulated image, and the
450 ground truth network, cumulative distribution functions were calculated for strand statistics
451 (length, average radius, average z-direction, and inverse tortuosity). Two comparisons were then
452 made for each strand statistics: Pearson's correlation between each upscaled or original
453 simulated image CDF and ground truth CDF; and Kolmogorov–Smirnov (K-S) test between each

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

upscaled image and original simulated image. The Pearson's correlation values compare the performance of each simulated (upscaled or original) image against the ground truth. The original simulated image serves as a baseline for vectorization performance. The K-S test is used to determine whether the performance of each model is significantly different compared to the original simulated image.

Overall accuracy was calculated for comparison as follows:

$$accuracy = 0.1(sensitivity) + 0.9(specificity)$$

with *sensitivity* and *specificity* defined as follows:

$$sensitivity = \frac{TP}{(TP + FN)}$$

$$specificity = \frac{TN}{(TN + FP)}$$

where TP is the number of true positive detections, TN is true negative, FP is false positive, and FN is false negative.

Blender renderings for vectorization visualization were created using VessMorphoVis software[20].

Statistical Analysis

All significance values were Bonferroni-adjusted from the standard P value of 0.05 to address the increased possibility of type-I error.

Data availability

All data used for training, validation, and testing are available at the following link: (link to dataverse)

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Code availability

All the code presented in this work together with the trained network model are freely available at the following link: (link to gitHub page)

References

1. Girouard, H. & Iadecola, C. Neurovascular coupling in the normal brain and in hypertension, stroke, and Alzheimer disease. *Journal of Applied Physiology* **100**, 328–335 (2006).
2. Lecrux, C. & Hamel, E. The neurovascular unit in brain function and disease: The neurovascular unit function. *Acta Physiologica* **203**, 47–59 (2011).
3. Meng, G. *et al.* Ultrafast two-photon fluorescence imaging of cerebral blood circulation in the mouse brain in vivo. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2117346119 (2022).
4. Hatakeyama, M., Ninomiya, I. & Kanazawa, M. Angiogenesis and neuronal remodeling after ischemic stroke. *Neural Regen Res* **15**, 16 (2020).
5. Recommendations for Standards Regarding Preclinical Neuroprotective and Restorative Drug Development. *Stroke* **30**, 2752–2758 (1999).
6. Shih, A. Y. *et al.* Two-Photon Microscopy as a Tool to Study Blood Flow and Neurovascular Coupling in the Rodent Brain. *J Cereb Blood Flow Metab* **32**, 1277–1309 (2012).
7. Voleti, V. *et al.* Real-time volumetric microscopy of in vivo dynamics and large-scale samples with SCAPE 2.0. *Nat Methods* **16**, 1054–1062 (2019).
8. Beaulieu, D. R., Davison, I. G., Kılıç, K., Bifano, T. G. & Mertz, J. Simultaneous multiplane imaging with reverberation two-photon microscopy. *Nat Methods* **17**, 283–286 (2020).
9. Mahou, P., Vermot, J., Beaurepaire, E. & Supatto, W. Multicolor two-photon light-sheet microscopy. *Nat Methods* **11**, 600–601 (2014).

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

10. Wu, J. *et al.* Kilohertz two-photon fluorescence microscopy imaging of neural activity in vivo. *Nat Methods* **17**, 287–290 (2020).
11. Sità, L. *et al.* A deep-learning approach for online cell identification and trace extraction in functional two-photon calcium imaging. *Nat Commun* **13**, 1529 (2022).
12. Soltanian-Zadeh, S., Sahingur, K., Blau, S., Gong, Y. & Farsiu, S. Fast and robust active neuron segmentation in two-photon calcium imaging using spatiotemporal deep learning. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 8554–8563 (2019).
13. Luo, L. *et al.* Restoration of Two-Photon Ca2+ Imaging Data Through Model Blind Spatiotemporal Filtering. *Front. Neurosci.* **15**, 630250 (2021).
14. Lee, S., Negishi, M., Urakubo, H., Kasai, H. & Ishii, S. Mu-net: Multi-scale U-net for two-photon microscopy image denoising and restoration. *Neural Networks* **125**, 92–103 (2020).
15. Tahir, W. *et al.* Anatomical Modeling of Brain Vasculature in Two-Photon Microscopy by Generalizable Deep Learning. *BME Frontiers* **2021**, 1–12 (2021).
16. Guan, H. *et al.* Deep-learning two-photon fiberscopy for video-rate brain imaging in freely-behaving mice. *Nat Commun* **13**, 1534 (2022).
17. Fang, L. *et al.* Deep learning-based point-scanning super-resolution imaging. *Nat Methods* **18**, 406–416 (2021).
18. Mihelic, S. A. *et al.* Segmentation-Less, Automated, Vascular Vectorization. *PLoS Comput Biol* **17**, e1009451 (2021).
19. Preibisch, S., Saalfeld, S. & Tomancak, P. Globally optimal stitching of tiled 3D microscopic image acquisitions. *Bioinformatics* **25**, 1463–1465 (2009).
20. Abdellah, M. *et al.* Interactive visualization and analysis of morphological skeletons of brain vasculature networks with VessMorphoVis. *Bioinformatics* **36**, i534–i541 (2020).

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

519 21. Zhou, A. *et al.* Evaluation of resonant scanning as a high-speed imaging technique for two-
520 photon imaging of cortical vasculature. *Biomed. Opt. Express* **13**, 1374 (2022).
521 22. Linninger, A., Hartung, G., Badr, S. & Morley, R. Mathematical synthesis of the cortical
522 circulation for the whole mouse brain-part I. theory and image integration. *Computers in*
523 *Biology and Medicine* **110**, 265–275 (2019).
524 23. Hartung, G. *et al.* Mathematical synthesis of the cortical circulation for the whole mouse
525 brain—part II: Microcirculatory closure. *Microcirculation* **28**, (2021).
526 24. Hopt, A. & Neher, E. Highly Nonlinear Photodamage in Two-Photon Fluorescence
527 Microscopy. *Biophysical Journal* **80**, 2029–2036 (2001).
528 25. Wang, T. *et al.* Quantitative analysis of 1300-nm three-photon calcium imaging in the mouse
529 brain. *eLife* **9**, e53205 (2020).
530 26. Cao, R. *et al.* Functional and oxygen-metabolic photoacoustic microscopy of the awake
531 mouse brain. *NeuroImage* **150**, 77–87 (2017).
532 27. Jafari, C. Z., Mihelic, S. A., Engelmann, S. & Dunn, A. K. High-resolution three-
533 dimensional blood flow tomography in the subdiffuse regime using laser speckle contrast
534 imaging. *J. Biomed. Opt.* **27**, (2022).
535 28. Schrandt, C. J., Kazmi, S. S., Jones, T. A. & Dunn, A. K. Chronic Monitoring of Vascular
536 Progression after Ischemic Stroke Using Multiexposure Speckle Imaging and Two-Photon
537 Fluorescence Microscopy. *J Cereb Blood Flow Metab* **35**, 933–942 (2015).
538 29. Engelmann, S. A. *et al.* Diamond Raman laser and Yb fiber amplifier for in vivo multiphoton
539 fluorescence microscopy. *Biomed. Opt. Express* **13**, 1888 (2022).
540 30. Schindelin, J. *et al.* Fiji: an open-source platform for biological-image analysis. *Nat Methods*
541 **9**, 676–682 (2012).

bioRxiv preprint doi: https://doi.org/10.1101/2022.11.30.518528; this version posted December 2, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Acknowledgments

V.M.N. was supported on a grant for human brain evolution by the Allen Discovery Center program, a Paul G. Allen Frontiers Group advised program of the Paul G. Allen Family Foundation as well as a fellowship from the Good Systems for Ethical AI at The University of Texas at Austin. This work was also supported by the National Institutes of Health (NS108484 to A.K.D; 3T32EB007507 and 5T32LM012414 to A.Z.). Objective and mouse images in Figure 1 were adapted from Biorender. GPU and computing support for the project was supported on Director's discretionary fund at the Texas Advanced Computing Cluster.

Contributions

A.Z. conceived of the study. A.Z., S.E., and A.T. collected data. A.Z., S.M., and V.M.N. analyzed the imaging data and built the pipeline. A.Z. and V.M.N. wrote the manuscript with input from all co-authors. V.M.N. and A.D. supervised the study.

Competing interests

The authors declare no competing interests.