Densely Populated Cell and Organelles Segmentation with U-Net Ensembles

Samuel J. Fulton[1], Cameron M. Baenen[1], Brian Storrie[2], Zijie Yang[1], Richard D. Leapman[1], and Maria A. Aronova[1]

[1]*National Institute of Biomedical Imaging and Bioengineering,*
National Institutes of Health, Bethesda, Maryland, USA and
[2]*Department of Physiology and Cell Biology, University of Arkansas for Medical Sciences, Little Rock, Arkansas, USA*

(Dated: January 7, 2025)

The complex and highly intertwined morphology of activated platelets within thrombi poses significant challenges for segmentation. In this work, we present a robust dual-network pipeline for cell and organelle segmentation. This multi-network approach enables the detection of fine details near the membrane while simultaneously facilitating long-range smoothing in regions distal to the membrane, drastically improving the performance of the watershed clustering algorithm compared to single-network approaches. We further enhance segmentation performance by collecting neural network predictions along multiple axes, capturing 3D correlations using only 2D neural networks. We segment and analyze hundreds of platelets and report quantitative morphological measurements, showing volumes consistent with hand-segmented results. We apply our segmentation pipeline to the CREMI neuron segmentation challenge data and provide state-of-the-art results.

INTRODUCTION

Recent advancements in high-throughput volume electron microscopy (vEM) have facilitated the acquisition of gigavoxel-scale datasets within a matter of hours. In stark contrast, manual segmentation of these datasets demands an extensive investment of human labor, often amounting to thousands of hours [1–3]. The emergence of deep learning methodologies has demonstrated remarkable efficacy in automating the traditionally labor-intensive process of segmentation [4–8]. Nonetheless, the segmentation of highly interconnected platelets within thrombi poses a particularly complex challenge. In this study, we employ lightweight neural networks in conjunction with geometric transformations to develop an advanced segmentation pipeline that achieves state-of-the-art performance.

Our segmentation pipeline employs a 2D U-Net architecture with a ResNet18 backbone, as illustrated in Fig. 1. The pipeline integrates two complementary networks: one trained to detect membranes and the other trained to estimate distances from membranes. By synthesizing the outputs of these networks, we generate a topographical landscape that is subsequently transformed into distinct cell instances using a watershed clustering algorithm. To enhance segmentation accuracy, we implement a multi-axis analysis approach, which aggregates network predictions across three orthogonal planes. This technique not only improves segmentation performance but also serves as a practical alternative to computationally intensive 3D neural networks. We evaluate our pipeline on both a platelet dataset and the CREMI Serial Block Face Scanning Electron Microscopy (SBF-SEM) neuron dataset [9], reporting metrics including precision, recall, F1 score, and accuracy.

RELATED WORK

Deep learning segmentation methodologies have advanced rapidly, encompassing supervised, semi-supervised, and unsupervised paradigms [2]. Fully supervised 3D segmentation techniques have achieved state-of-the-art performance, enabling cellular analysis at scales previously deemed infeasible [10–12]. Similarly, weakly supervised and unsupervised techniques have demonstrated success in scenarios characterized by sparse data availability [13–15].

The development of this field has been significantly supported by publicly available FIB-SEM datasets, such as Lucchi++ (mouse hippocampus), FIB-25 (drosophila optic lobe), OpenOrganelle (Hela cells), Cardiac Mitochondria (mouse heart), and UroCell (mouse urothelial cells) [16]. Leveraging these resources, our work extends the application of deep learning techniques to a novel dataset: blood platelets within mouse thrombi.

The conventional approach to instance-level segmentation typically employs a U-Net architecture to generate probability maps, which indicate the likelihood of a pixel belonging to a specific object. These maps are subsequently processed using classical algorithms, such as watershed clustering, to achieve instance-level segmentation [14, 17–21]. However, pixel classification methods often suffer from the erroneous splitting of large cells. To mitigate this limitation, alternative approaches have been developed, wherein neural networks are trained to predict distance metrics from the background. Euclidean distance-based methods have proven effective in preventing the fragmentation of large cells [14, 22–27]. Nevertheless, experimental observations reveal that these distance-mapping techniques can lead to the misidentification of dendrites and other branching structures at cell surfaces. To address these challenges, we integrate standard membrane identification with distance mapping, combining the strengths of both methodologies.

To further enhance segmentation accuracy, we exploit the volumetric information inherent in the data. While

volumetric approaches in most vEM datasets are often hindered by lower z-resolution, necessitating cross-slice attention models that emulate the decision-making process of a radiologist [12, 28–30], the isotropic nature of FIB-SEM datasets enables a more elegant solution. By rotating the data to collect neural network predictions along multiple axes, we efficiently capture 3D correlations using only a 2D neural network [27, 31].

Our work advances platelet segmentation by transitioning from semantic segmentation, as demonstrated by Guay et al. [1], to instance-level segmentation. By focusing on platelets within thrombi, we extend the research of Matharu et al., which analyzed platelets in cell culture [20]. To the best of our knowledge, this represents the first large-scale segmentation of platelets within thrombi, offering a significant contribution to the field.

PLATELET DATASET AND MULTI-AXIS ANALYSIS

We center the analysis on a FIB-SEM platelet dataset from a mouse thrombus, which is comprised of tightly arranged blood platelets. The FIB-SEM dataset is stored in a 3D tiff stack with dimensions $2796 \times 2388 \times 1012$ and near isotropic voxel size of $5 \times 5 \times 7$ nm. We perform instance-level segmentation of mitochondria, α-granules, and whole platelets.

To address the limitations of conventional 2D segmentation approaches, we employ a 2D neural network applied along three orthogonal axes. Initially, the network is executed on each 2D slice of the TIFF stack. The stack is then rotated, and the network is applied to the new forward-facing plane. This process is repeated a second time, ensuring that the network processes the data along three orthogonal axes. The resulting predictions are integrated into a unified 3D probability map, a technique we term multi-plane analysis. This approach is grounded in the geometric principle that membranes parallel to the forward-facing plane are often indistinguishable from the background, leading to potential false negatives in U-Net predictions. Since a membrane can only be parallel to one of the three orthogonal planes, multi-axis analysis significantly reduces the likelihood of false negatives, confining them to a single plane.

The use of 2D neural networks with multi-plane analysis offers several advantages over 3D neural networks while retaining much of their predictive capability. First, generating 3D training datasets is considerably more challenging compared to 2D datasets. Second, the availability of pre-trained 2D backbones, such as ResNet18 and VGG16, facilitates faster training and reduces the dependency on extensive manually labeled datasets [32]. This approach thus provides a computationally efficient and resource-effective alternative to 3D neural networks, without compromising segmentation performance.

FIG. 1: (A) Hand-labeled platelet membranes used for creating distance maps as well as training the neural network. (B) Grey-scale image showing the distance from the platelet membrane. (C) U-Net architecture with ResNet18 pre-trained backbone and decoder layers, utilizing three residual connections. Network input is a greyscale image and output is a probability or distance map.

NEURAL NETWORK ARCHITECTURE

Our network architecture is based on the U-Net framework, initially introduced by Olaf Ronneberger et al. in 2015 [33, 34]. Since its inception, the U-Net has been extensively employed in biomedical imaging [5–7]. The encoder of our U-Net consists of layers from ResNet18, which are split into a convolution layer and four subsequent ResNet18 blocks, as shown in 1. The bottleneck consists of two convolution layers and ReLU activations. The decoder consists of three convolution and up-sampling blocks. We utilize three residual connections, which help stabilize the training process. Lastly, a 1×1 convolution layer with sigmoid activation produces the segmentation probability and distance maps.

U-NET TRAINING

We employ a single-channel neural network. A is motivated by the lack of overlapping hand-labeled cell membrane and organelle data, making the single-channel approach more practical for our purposes. The Whole-Cell Segmentation Pipeline is based on two neural networks trained to complementary representations of our data.

The training dataset was obtained by hand-labeling membranes using Amira Avizo commercial software. A total of fourteen randomly selected images each 2796×2388 pixels were labeled. After labeling the membranes we used the Amira Avizo distance "mapping" module

to create the corresponding distance maps. We divided our labeled data into 90 percent training and 10 percent validation data. We augmented the training data by applying elastic deformations, rotations, brightness shifts, contrast shifts, and Gaussian noise.

We randomly cropped each image into 100 patches of 1024×1024 pixels with each patch receiving a random combination of augmentations. We train for twenty epochs with a batch size of 32 and use early dropout with a patience of five epochs. We utilize Adam optimizer with learning rate of 0.00001 and L2 regularization. We train using a single A100 GPU on NIH's computing cluster (Biowulf), where training required thirty minutes [35]. We also train using a local consumer-grade GeForce GTX 1080 GPU which required a longer training time of two hours.

For generating the membrane probability maps we use binary cross entropy as the loss function. For the training of the distance mapping, we define a custom loss function.

$$\text{Loss} = \frac{(\text{predictions} - \text{targets})^2}{\log(2 + \text{targets})},$$

which more heavily weighs the error closer to the membrane. The loss function is given by the mean squared error scaled to be large near the membrane and smaller away from the membrane. By weighting the loss more heavily near the membrane, we achieve higher detail near the borders of cells, preventing cells from incorrectly merging.

SEGMENTATION PIPELINES

To convert the neural network outputs to cell and organelle instances we use Amira Avizo's watershed segmentation tools [36]. The watershed segmentation technique interprets pixel intensities as a topographic surface, "flooding" valleys from predefined markers to segment regions based on their topographical boundaries [37]. Numerous software packages and products offer implementations of the watershed algorithm; however, despite sharing a common conceptual basis, their performance can differ significantly [38]. Amira Avizo, for instance, features three primary watershed implementations: H-Extrema, Hierarchical, and Marker-Based, each suited for unique topographical landscapes. Due to differences in the landscape of the organelle and cell membrane probability maps, we propose two segmentation pipelines based on different watershed techniques. Both segmentation pipelines are outlined in Fig.2

Whole-Cell Segmentation Pipeline

The whole-cell segmentation pipeline is a multi-step process illustrated in Fig.2. First, we use two U-Nets to identify cell membranes and distances from membranes

and utilize multi-axis analysis to capture volumetric information. We normalize distances and probability maps to 1 and add the results. We use H-Extrema watershed to convert the combined probability/distance map into cell instances. We employ the H-extrema watershed method over other watershed techniques because it offers a single, tunable parameter that defines the minimum distance between neighboring cells. For tightly packed platelets, this distance is nearly uniform, making optimal parameter values easy to find. The H-extrema watershed algorithm uses regions of low cell membrane likelihood as initial seeds, expanding outward until adjacent cells meet. Cells located within the distance defined by the tunable parameter are merged, ensuring effective segmentation in densely packed cellular structures.

Organelle Segmentation Pipeline

To segment organelles, we employ a dual network pipeline. We train one neural network to identify the organelle membrane and one neural network to identify the entire organelle. We run both neural networks on the data, creating two sets of probability maps. We subtract the membrane map from the whole organelle map to create a probability map of the organelle contents. We use hierarchical watershed to segment only organelle contents (rather than full organelle) ensuring that neighboring organelles are not merged. We then use marker-based watershed expanding the organelles to include their membranes.

BENCHMARKING ON CREMI DATASET

We benchmark the whole-cell segmentation pipeline on CREMI neuron segmentation challenge Sample_A. The CREMI dataset comprised 125 hand-labeled slices each with a dimension of 1250×1250 pixels. Unlike our FIB-SEM dataset with a near-isotropic voxel shape, the CREMI datasets have a voxel shape of $4 \times 4 \times 40$ nm. To utilize multi-plane analysis, we train two neural networks. The first neural network is trained using 12 slices from the dataset and can be run in the xy-plane. We create the training dataset for the second neural network by resampling the 12 slices into 120 slices each with pixel shape 4×40 such that the pixel shape matches that of the xz and yz-planes. A slice of our segmentation results is shown next to the ground truth segmentation in Fig.3.

To benchmark the final segmentation output, we assign pixels near the cell instance boundary to the membrane and compare them with the ground truth membrane. Fig.3 shows ground truth overlaid with membranes from the whole-cell segmentation pipeline. We calculate precision, recall, accuracy, and F1 scores and show results in Table.I. Our pipeline shows quality results, classifying pixels correctly with a precision of 0.923. Some difficulties arise in running watershed on anisotropic datasets.

A: Whole-Cell Segmentation Pipeline

B: Organelle Segmentation Pipeline

FIG. 2: We propose two complementary pipelines for cell and organelle segmentation. Both pipelines are illustrated with platelet data. The **Whole-Cell Segmentation Pipeline** performs well with cases of uniform cell density. **(A.1)** shows the raw data. **(A.2)** shows membrane probability map averaged from multi-axis analysis. Similarly, **(A.3)**, shows the predicted distance to the nearest membrane averaged from multi-axis analysis. **(A.4)** shows final cell instances created using H-extrema watershed. The **Organelle Segmentation Pipeline** performs much better for cases of varying instance density. We employ two neural networks: one to segment the entire organelle **(B.2)** and one to segment organelle membrane **(B.3)**. We subtract the membrane from the full organelle results, **(B.4)**, and use Hierarchical watershed to create instance segmentation of the interior of the organelles **(B.5)**. We then used the interiors and the full organelles probability mask to perform marker-based watersheds, creating instances of full organelles **(B.6)**.

FIG. 3: **(A)** Raw FIB-SEM data. **(B)** Whole cell segmentation results. **(C)** Raw SBF-SEM neuron data from CREMI segmentation challenge. **(D)** Whole cell segmentation results.

mitochondria, and cell membranes. While instance-level segmentations are generated for α-granules, mitochondria, and platelets, these classes are converted to semantic labels for benchmarking purposes. Our results demonstrate state-of-the-art accuracy and precision, consistently exceeding 0.85. Notably, the membrane detection pipeline performs exceptionally well, a critical requirement for accurate morphological measurements. In contrast, slightly lower performance is observed for the mitochondria and α-granule benchmarks, primarily due to missed pixels near organelle edges. However, key metrics of organelle distribution, which are closely associated with disease, remain robust to such errors [39, 40].

Metric	Mitochondria	Alpha Granules	Platelet Membrane
Precision	0.96	0.91	0.99
Recall	0.60	0.79	0.96
F1 Score	0.74	0.85	0.98
Accuracy	0.99	0.97	0.99

TABLE II: The metric formulas are defined in Table.I. Membrane classification demonstrates state-of-the-art performances in all metrics.

While the U-Net utilizes the high xy-resolution, the watershed algorithm suffers from low z-resolution. Andong Wang et al. implemented an algorithm known as Touching Area-Based Clustering Algorithm (TASCAN) to improve watershed results on anisotropic data [39]. We do not implement TASCAN here. Even with the difficulties with running watershed on anisotropic data, we still see good membrane classification results with accuracy and precision over 0.85.

Metric	Formula	Neuron Membrane
Precision	$\frac{TP}{TP+FP}$	0.923
Recall	$\frac{TP}{TP+FN}$	0.756
Accuracy	$2\frac{\text{Precision}\times\text{Recall}}{\text{Precision}+\text{Recall}}$	0.875
F1 Score	$\frac{TP+TN}{TP+TN+FP+FN}$	0.831

TABLE I: The metric formulas are expressed in terms of True Positive (TP), True Negative (TN), False Positive (FP), and False Negative (FN).

BENCHMARKING PLATELET SEGMENTATION RESULTS

We present segmentation results for the FIB-SEM platelet dataset, targeting the detection of α-granules,

3D RENDERINGS AND MORPHOLOGICAL MEASUREMENTS

We use Amira Avizo to display segmentation results of mitochondria, α-granules, and cell instances in Fig.4. The 3d renderings were created using Amira Avizo's "Generate Surface" and "Surface View" modules. We apply Amira Avizo's "Label-Analysis" tool to calculate volumes and axis lengths for all segmented objects and export the results as CSV files to be analyzed with Excel and Python. We report several parameters of interest and compare them with hand-labeled results published in Table.III. Our measured cell volume and published results, 2.69 ± 0.12 (μm^3) and 2.90 ± 0.34 (μm^3), are within a single error bar of each other, consistent with the understanding that activation will affect platelet shape but not mass. Other morphological measurements differ due to differences in environment and platelet activation levels. The 3D rendering and morphological measurements provide insight into organelle clustering, cellular interactions, and cell polarity. We hope that large-scale segmentation such as the one performed here will help provide insight into the structural underpinnings of blood clot formation and heart disease.

We also segment platelets from four different regions of the clot and show 3d-rendering in Fig.5. The segmentation pipeline readily generalizes to multiple FIB-SEM datasets.

FIG. 4: (A) FIB-SEM data overlaid with corresponding mitochondria instance segmentation results. (B) FIB-SEM data overlaid with α-granule instance segmentation results. (C) FIB-SEM data overlaid with whole-cell segmentation results. (D) 3D rendering showing 127 cells from the interior of a thrombus. (E) 3D rendering showing a single cell from the center of a thrombi.

FIG. 5: Cells from regions of interest. **A** Shows cells near the vessel wall. **B** Shows cells near a gap in the thrombus. C Shows cells near the center of the thrombus. **D** shows cells near the vessel wall.

OUTLOOK

Here we present the lightweight pipeline for cell and organelle segmentation and morphological analysis. We readily apply the pipeline to the segmentation of cells within thrombi. We lay a foundation for future efforts to improve automated segmentation in biomedical contexts. We are currently analyzing numerous platelet datasets to compare ultrastructure in differing regions of a thrombus.

Parameter	Our Results $(\mu \pm \sigma/\sqrt{n})$	Published $(\mu \pm \sigma/\sqrt{n})$
Platelet Volume (μm^3)	2.69 ± 0.12	2.90 ± 0.34
Platelet Minor and Major Axis (µm)	1.74 ± 0.62 3.74 ± 1.67	0.87 ± 0.07 2.70 ± 0.15
Alpha-Granule Number/Platelet	11 ± 1	19 ± 4.3
Alpha Granule Short and Long Axis (µm)	0.364 ± 0.022 0.464 ± 0.005	0.288 ± 0.006 0.537 ± 0.018
Volume/Alpha Granule (µm³)	0.025 ± 0.00061	0.013 ± 0.00076
Alpha Granule Volume Fraction/Platelet	0.102 ± 0.005	0.088 ± 0.007
Mitochondria Number/Platelet	5.10 ± 0.29	12 ± 1.8
Volume/Mitochondria (µm³)	0.012 ± 0.00027	0.0072 ± 0.00015
Mitochondria Volume Fraction/Platelet	0.023 ± 0.003	0.029 ± 0.002

TABLE III: Measured ultrastructure parameters of interest compared with previously published data. Our calculated volume shows good agreement with published results. Other parameters differ due to activation levels. Published data retrieved from [41].

This pipeline will provide unique insight into thrombi formation and heart disease. Furthermore, by enabling extensive quantification of large datasets, our framework provides a basis for robust research into cellular dynamics, activation, and environments. We design our pipeline for accessibility and hope to bridge the gap between advanced machine-learning methods and practical applications in experimental biology.

ACKNOWLEDGMENTS

We thank Joshua Kim for collecting the FIB-SEM data, Elizabeth R. Driehaus and Sidney W. Whiteheart for performing the surgeries, Irina D. Pokrovskaya for sample prep, Guofeng Zhang for sample embedding, and the Storrie lab for insightful biological discussions motivating this project.

This work is supported by the intramural research program of the National Institute of Biomedical Imaging and Bioengineering. Research in the Storrie laboratory is supported in part by NIH grant R01 HL155519. This work utilized the computational resources of the NIH HPC Biowulf cluster (`http://hpc.nih.gov`).

CODE AVAILABILITY

All code can be found at our public GitHub repository: `https://github.com/fultonsaj/Deep-Learning-Cell-Seg/tree/main`

[1] M. D. Guay, Z. A. S. Emam, A. B. Anderson, M. A. Aronova, I. D. Pokrovskaya, B. Storrie, and R. D. Leapman, Dense cellular segmentation for em using 2d–3d neural network ensembles, Scientific Reports https://doi.org/10.1038/s41598-021-81590-0 (2021).

[2] A. Aswath, A. Alsahaf, B. N. Giepmans, and G. Azzopardi, Segmentation in large-scale cellular electron microscopy with deep learning: A literature survey, Medical Image Analysis (2023).

[3] M. Berning, K. M. Boergens, and M. Helmstaedter, Segem: efficient image analysis for high-resolution connectomics, Neuron **87**, 1193 (2015).

[4] M. Helmstaedter, K. Briggman, and W. Denk, Connectomic reconstruction of the inner plexiform layer in the mouse retina, Nature **500**, 168 (2013).

[5] Z. Zhou, M. M. R. Siddiquee, N. Tajbakhsh, and J. Liang, Unet++: A nested u-net architecture for medical image segmentation, in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support* (Springer, 2018) pp. 3–11.

[6] J. Schlemper, O. Oktay, M. Schaap, M. Heinrich, B. Kainz, B. Glocker, and D. Rueckert, Attention gated networks: Learning to leverage salient regions in medical images, Medical Image Analysis **53**, 197 (2019).

[7] N. Ibtehaz and M. S. Rahman, Multiresunet: Rethinking the u-net architecture for multimodal biomedical image segmentation, Neural Networks **121**, 74 (2020).

[8] S. M. Plaza and J. Funke, Analyzing image segmentation for connectomics, Frontiers in Neural Circuits **12**, 102 (2018).

[9] P. Li, J. Funke, Ariadne, C. Ordish, S. Saalfeld, and J. Funke, Miccai challenge on circuit reconstruction from electron microscopy images (2017).

[10] L. Heinrich, D. Bennett, D. Ackerman, W. Park, J. Bogovic, N. Eckstein, A. Petruncio, J. Clements, S. Pang, C. Xu, J. Funke, W. Korff, H. Hess, J. Lippincott-Schwartz, S. Saalfeld, and A. Weigel, Whole-cell organelle segmentation in volume electron microscopy, Nature (2021).

[11] A. Shapson-Coe, M. Januszewski, D. Berger, A. Pope, Y. Wu, T. Blakely, R. Schalek, P. Li, S. Wang, and J. Maitin-Shepard, A petavoxel fragment of human cerebral cortex reconstructed at nanoscale resolution, Science **384** (2024).

[12] H. Jia, Y. Xia, Y. Song, D. Zhang, H. Huang, Y. Zhang, and W. Cai, 3d apa-net: 3d adversarial pyramid anisotropic convolutional network for prostate segmentation in mr images, IEEE Transactions on Medical Imaging **39**, 447 (2019).

[13] A. Wolny, Q. Yu, C. Pape, and A. Kreshuk, Sparse object-level supervision for instance segmentation with pixel embeddings, in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (IEEE, Heidelberg, Germany, 2022).

[14] J. Ma, Z. Wei, Y. Zhang, Y. Wang, R. Lv, C. Zhu, C. Gaoxiang, J. Liu, C. Peng, L. Wang, Y. Wang, and J. Chen, How distance transform maps boost segmentation cnns: An empirical study, in *Proceedings of the Third Conference on Medical Imaging with Deep Learning*, Proceedings of Machine Learning Research, Vol. 121, edited by T. Arbel, I. B. Ayed, M. de Bruijne, M. Descoteaux, H. Lombaert, and C. Pal (PMLR, 2020) pp. 479–492.

[15] R. Bermúdez-Chacón, O. Altingövde, C. Becker, M. Salzmann, and P. Fua, Visual correspondences for unsupervised domain adaptation on electron microscopy images, IEEE Transactions on Medical Imaging **39**, 1256 (2020).

[16] A. Aswath, A. Alsahaf, B. Giepmans, and G. Azzopardi, Segmentation in large-scale cellular electron microscopy with deep learning: a literature survey, Medical Image Analysis **89**, 102920 (2023).

[17] D. Ciresan, A. Giusti, L. M. Gambardella, and J. Schmidhuber, Deep neural networks segment neuronal membranes in electron microscopy images, in *Advances in Neural Information Processing Systems 25*, edited by F. Pereira *et al.* (Neural Information Processing Systems Foundation, 2012) pp. 2852–2860.

[18] S. C. Turaga, J. F. Murray, V. Jain, F. Roth, M. Helmstaedter, K. Briggman, W. Denk, and H. S. Seung, Convolutional networks can learn to generate affinity graphs for image segmentation, Neural Computation **22**, 511 (2010).

[19] W. Shen, B. Wang, Y. Jiang, Y. Wang, and A. Yuille, Multi-stage multi-recursive-input fully convolutional networks for neuronal boundary detection, in *2017 IEEE International Conference on Computer Vision (ICCV)* (IEEE, Venice, Italy, 2017) pp. 2410–2419.

[20] S. S. Matharu, C. S. Nordmann, K. R. Ottman, R. Akkem, D. Palumbo, D. R. D. Cruz, R. D. Leapman, and B. Storrie, Deep learning, 3d ultrastructural analysis reveals quantitative differences in platelet and organelle packing in covid-19/sarscov2 patient-derived platelets, Platelets **34**, 2264978 (2023).

[21] D. Franco-Barranco, A. Muñoz-Barrutia, and I. Arganda-Carreras, Stable deep neural network architectures for mitochondria segmentation on electron microscopy volumes, Neuroinformatics **20**, 437 (2022).

[22] W. Wang, D. A. Taft, Y.-J. Chen, J. Zhang, C. T. Wallace, M. Xu, S. C. Watkins, and J. Xing, Learn to segment single cells with deep distance estimator and deep cell detector, Computers in Biology and Medicine **108**, 133 (2019).

[23] U. Schmidt, M. Weigert, C. Broaddus, and G. Myers, Cell detection with star-convex polygons. medical image computing and computer assisted intervention—miccai 2018 (2018).

[24] S. E. A. Raza, L. Cheung, M. Shaban, S. Graham, D. Epstein, S. Pelengaris, M. Khan, and N. M. Rajpoot, Micronet: A unified model for segmentation of various objects in microscopy images, Medical image analysis **52**, 160 (2019).

[25] X. Zhou, F. Li, J. Yan, and S. T. C. Wong, A novel cell segmentation method and cell phase identification using markov model, IEEE Transactions on Information Technology in Biomedicine **13**, 152 (2009).

[26] P. Naylor, M. Laé, F. Reyal, and T. Walter, Segmentation of nuclei in histopathology images by deep regression of the distance map, IEEE transactions on medical imaging **38**, 448 (2018).

[27] H. Spiers, H. Songhurst, L. Nightingale, J. de Folterand Christopher J Peddie, A. Weston, A. Strange, S. Hindmarsh, C. Lintott, L. M. Collinson, and M. L. Jones, Deep learning for automatic segmentation of the nuclear envelope in electron microscopy data, trained with volunteer segmentations traffic, John Wiley Online Library 10.1111/tra.12789 (2021).

[28] Y. Al-Kofahi, A. Zaltsman, R. Graves, W. Marshall, and M. Rusu, A deep learning-based algorithm for 2-d cell segmentation in microscopy images, BMC Bioinformatics **19**, 365 (2018).

[29] H. A. Yu, H. Zheng, K. Zhao, X. Du, K. Pang, Q. Miao, S. S. Raman, D. Terzopoulos, and K. Sung, Csam: A 2.5d cross-slice attention module for anisotropic volumetric medical image segmentation, in *IEEE Winter Conference on Applications of Computer Vision (WACV)* (2024) pp. 5911–5920.

[30] Z. Huang, Y. Zhao, Y. Liu, and G. Song, Gcaunet: A group cross-channel attention residual unet for slice-based brain tumor segmentation, Biomedical Signal Processing and Control **70**, 102958 (2021).

[31] C. Stringer, T. Wang, M. Michaelos, and M. Pachitariu, Cellpose: a generalist algorithm for cellular segmentation, Nature Methods https://doi.org/10.1101/2020.02.02.931238 (2021).

[32] H. Kaiming, Z. Xiangyu, R. Shaoqing, and S. Jian, Deep residual learning for image recognition, in *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)* (2016).

[33] N. Siddique, S. Paheding, C. P. Elkin, and V. Devabhaktuni, *U-Net and Its Variants for Medical Image Segmentation: A Review of Theory and Applications* (IEEE, 2021).

[34] O. Ronneberger, P. Fischer, and T. Brox, U-net: Convolutional networks for biomedical image segmentation, Springer Cham https://doi.org/10.1007/978-3-319-24574-4 (2015).

[35] National Institutes of Health, Biowulf high-performance computing, `https://hpc.nih.gov` (2024), accessed: 2024.

[36] Thermo Fisher Scientific, Amira-Avizo Software for 3D Visualization and Analysis (2024), [Software]. Available from: `https://www.thermofisher.com`.

[37] B. Serge and L. Christian, Use of watersheds in contour detection, In International Workshop on Image Processing: Real-time Edge and Motion Detection/Estimation (1979).

[38] A. Kornilov, I. Safonov, and I. Yakimchuk, A review of watershed implementations for segmentation of volumetric images, Journal of Imaging https://doi.org/10.3390/jimaging8050127 (2022).

[39] N. Nešić, X. Heiligenstein, L. Zopf, V. Blüml, K. S. Keuenhof, M. Wagner, J. L. Höög, H. Qi, Z. Li, G. Tsaramirsis, C. J. Peddie, M. Stojmenović, and A. Walter, Automated segmentation of cell organelles in volume electron microscopy using deep learning, Microscopy Research and Technique https://doi.org/10.1002/jemt.24548 (2024).

[40] D. F. Silva, J. E. Selfridge, J. Lu, L, E, S. M. Cardoso, and R. H. Swerdlow, Mitochondrial abnormalities in alzheimer's disease: possible targets for therapeutic intervention, Advances in Pharmacology **64**, 83 (2012).

[41] I. Pokrovskaya, M. Tobin, R. Desai, M. A. Aronova, J. A. Kamykowski, G. Zhang, S. Joshi, S. W. White-

heart, R. D. Leapman, and B. Storrie, Structural analysis of resting mouse platelets by 3d-em reveals an unexpected variation in alpha-granule shape, Platelets 10.1080/09537104.2020.1799970 (2021).