Deep Learning Segmentation and Analysis of Neural Organoids for Schizophrenia Research

Jack Bodine *University of Copenhagen* qnj501@alumni.ku.dk
Lucrezia Criscuolo *University of Southern Denmark* lucr@sdu.dk
Barbara Lykke Lind *University of Copenhagen* barbarall@sund.ku.dk
Svetlana Kutuzova *University of Copenhagen* svegal@di.ku.dk

Abstract

Manual region-of-interest annotation in two-photon calcium imaging is a necessary step for meaningful results, yet remains a tedious process, limiting scalable studies of neural organoids. We built an AI pipeline using U-Net to segment neural organoids and quantify activity, enabling cell-level investigations into schizophrenic and control organoids. Using just 12 ground truth images, we trained a U-Net segmentation model that achieved a macro Dice = 0.8321 on the evaluation set. We used these masks to measure cell specific activity. The activity traces were then used to analyze differences in calcium activity between schizophrenic and control subsets. The pipeline demonstrates the ability AI has to accelerate neuroscience research.

Keywords: Image segmentation, Neural organoids, Microscopy, Calcium imaging, Time-series analysis

1 Introduction

Neural organoids coupled with two-photon calcium imaging (2PCI) is an emerging system for researching the human brain and gaining an understanding of neurological diseases at a cellular level (Grienberger et al. [2022]). This approach currently requires hours of manual annotation of data by trained neuroscience experts and produces large volumes of output that is challenging to sift through. Existing 2PCI segmentation tools for in vivo neurons often underperform when applied to the irregular shapes and sizes of neural organoids. Using U-Net for cell-level segmentation of organoids and feature engineering, we create a new deep-learning pipeline that can reveal fundamental insights into the human brain.

The neural organoids analyzed in this paper are derived from individuals with schizophrenia (SCZ) and healthy controls (CTL). Our AI pipeline makes it possible to compare the two groups, which can reveal biological differences and advance our understanding of the disease.

2 Background

2.1 Neural Organoids

Neural organoids are lab-grown, three-dimensional in vitro tissues generated from patient-derived induced pluripotent stem cells and composed of diverse cell populations. They provide neuroscientists a non-invasive method for studying neurons in human tissue (Lee and Sun [2022]). Similar to in vivo neurons, they have measurable amounts of fluctuations in the concentration of calcium ions (Ca^{2+}), referred to as calcium activity. Using two-photon calcium imaging (2PCI), we infer this activity from changes in fluorescence of calcium indicators (Grienberger et al. [2022]), tracking brightness over time to estimate each neuron's activity.

2.2 Segmentation Models

Deep learning approaches have achieved state-of-the-art results in biomedical imaging tasks across many different applications. U-Net (Ronneberger et al. [2015]) is one of the top performing segmentation models. Deep convolutional neural networks (CNNs), the building blocks of U-Net, are also able to perform segmentation, but usually require large datasets. The U-Net model architecture overcomes this issue.

3 Methodology

3.1 Data

The dataset used throughout this paper was composed of 693 2PCI videos of organoids. In each video, the amount of calcium activity occurring in specific neuronal cell types was recorded in the green channel, while calcium activity in all neurons was recorded in the red channel.

The samples in this dataset were derived from six individuals, three with schizophrenia and three healthy controls. For each cell line, three biological replicates were analyzed for each of the two neuronal cell types: glutamatergic and GABAergic. Every subsample had between one and seven field-of-views (FOVs) recorded, for a total of 225 FOVs. Lastly, each FOV had between 3 and 5 recordings, totaling 693 files.

We created mean-projection images from each 2PCI video to serve as inputs for segmentation. Example mean projection images are shown in Figure 1. For 22 of the mean-projection images, we created ground truth region of interest (ROI) segmentation masks by labeling where each cell is located. We included both the cell body, called the soma, and its axons and dendrites, which together are referred to as processes.

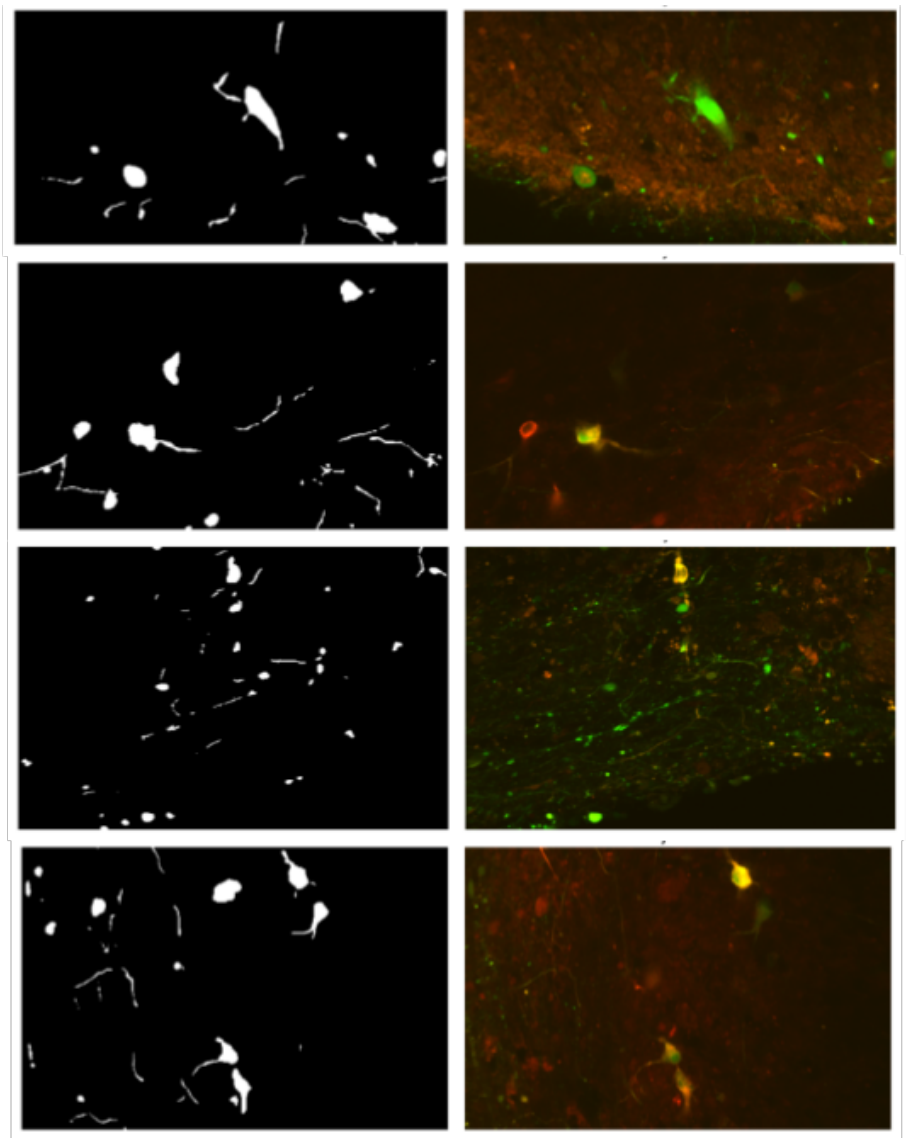

Figure 1: Four selected 2PCI mean projection images (right) and their corresponding ROI segmentation masks (left) predicted by the U-Net trained on all annotated samples.

3.2 Generating Region of Interest Masks with the U-Net Architecture

We trained U-Net segmentation models using Yucca, a PyTorch framework built for medical imaging tasks (Llambias et al. [2024]). Yucca provided modular components for data I/O, augmentations, and standardized training/inference. Its multi-modality support let us feed both the red (all neurons) and green (targeted neurons) channels as separate inputs while sharing a single ground truth. Organoids remain stationary in the dish, so no motion correction was required.

To mitigate limited training data, we applied three types of augmentations: rotations, random crops, and intensity fluctuations. Each augmentation type individually improved performance, with rotations contributing the largest gains. Each image had 35 rotations made at 10-degree intervals, each rotation had 8 random crops applied, each crop had 4 intensity modifications.

We trained on the augmented images derived from 12 ground truth images and evaluated on 10 held-out images, obtaining a macro Dice coefficient of 0.8321 on the test set. Figure 1 shows example inputs and the corresponding U-Net generated ROI masks. The model accurately localizes most somas and processes. In some cases, the soma or process being segmented has very small or difficult to trace details which would likely be missed by a human annotator—our model often includes these objects even if they are only a couple pixels in size. Some somas remain excluded from the segmentation mask, particularly those of low intensity which only appear in the red channel. This could possibly be due to these objects being underrepresented in the training set.

The model presented here outperformed various other open source options. One open source model, *Cell Identification and Trace Extraction Online* (CITE-on), is a convolutional neural network model for segmenting neural cells in vivo and deriving their activity traces from 2PCI data (Sità et al. [2022]). However, when applied to the neural organoid dataset, CITE-on failed to correctly identify most of the cells. The in vivo neurons CITE-on is trained to identify have a much higher density, more uniform size, and slightly different appearance than the organoids in our dataset. Additionally, CITE-on only locates and measures soma activity while our model is trained to identify both somas and processes (axons and dendrites).

The U-Net model we trained also outperformed a U-Net using a pretrained ResNet model as its encoder, referred to as a "U-Net with a ResNet backbone.' We evaluated this model using three ResNet models (50, 101, 152) each pretrained on the ImageNet dataset. The best performing non-fine-tuned model (ResNet50) achieved a macro F1-score of 0.388, with fine-tuning the best performing model (ResNet152) achieved a macro F1-score of 0.742, still underperforming against our model.

Given these results, we proceeded to use our U-net model's masks for downstream cell-based activity analysis. We trained a production model on all 22 labeled images and used it to generate masks for the remaining videos.

3.3 Post Mask Analysis

Using the U-Net ROI masks, we extracted per-cell calcium activity traces from each 2PCI video by summing pixel intensities under each mask for each frame. This gives us a per-cell signal that tracks activity over time. Each object in the masks was classified as a soma or processes using a roundness threshold (0.65) and size threshold (8 µm).

From these calcium activity signals we calculated time, frequency domain, and wavelet features for each cell and for the sample as a whole. We also calculated neuron specific features e.g., burst rate and mean spikes per burst. We then analyzed these features for differences between the schizophrenic and control groups.

We found several group differences with statistical significance ($p < 0.05$). Including spectral flatness, event rate, cell intensities, and process density. At this stage, these findings are exploratory.

4 Conclusion and Future Work

Our work shows that AI tools can be used to unlock discoveries in neural organoids and 2PCI data. This work delivers an AI-backed pipeline that replaces manual ROI annotation with learned segmentation and analysis for 2PCI neural organoids, enabling scalable SCZ/CTL comparisons. We have produced a domain-specific U-Net with targeted augmentations that achieved a macro Dice coefficient of 0.8321 on just 12 labeled images, then produced masks and features for the remaining data. This pipeline accelerates joint work between neuroscience and computer science and provides a transferable foundation for new datasets and future work.

Our next steps are to determine whether these models and this pipeline can be applied to different datasets. We are currently investigating whether the U-Nets here can transfer to 2PCI data collected from in vivo neurons from mice. If they do, we plan to use this pipeline to study the mechanisms of ADHD treatments in naturally exploring mice. Additionally, we are curious if we can use larger patient-level organoid cohorts to train feature-based classifiers that can predict schizophrenia from an organoid's activity profile.

References

C. Grienberger, A. Giovannucci, W. Zeiger, and C. Portera-Cailliau. Two-photon calcium imaging of neuronal activity. *Nature Reviews Methods Primers*, 2(1), Sep 2022. doi: 10.1038/s43586-022-00147-1.

J.-H. Lee and W. Sun. Neural organoids, a versatile model for neuroscience. *Molecules and Cells*, 45(2):53–64, Feb 2022. doi: 10.14348/molcells.2022.2019.

S. N. Llambias, J. Machnio, A. Munk, J. Ambsdorf, M. Nielsen, and M. M. Ghazi. Yucca: A deep learning framework for medical image analysis, 2024. URL `https://arxiv.org/abs/2407.19888`.

O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation, 2015. URL `https://arxiv.org/abs/1505.04597`.

L. Sità, M. Brondi, P. Lagomarsino de Leon Roig, S. Curreli, M. Panniello, D. Vecchia, and T. Fellin. A deep-learning approach for online cell identification and trace extraction in functional two-photon calcium imaging. *Nature Communications*, 13(1), 3 2022. ISSN 2041-1723. doi: 10.1038/s41467-022-29180-0. URL `http://dx.doi.org/10.1038/s41467-022-29180-0`.