EE367: Final Project Proposal

Benny Weng, Alex Kiral

February 20, 2026

1 Motivation

Two-photon microscopy is considered a workhorse in modern neuroscience due to its superior optical sectioning capability (subcellular resolution both laterally and axially) and increased scattering length resulting from the utilization of near infrared (NIR) excitation sources. The application of two-photon excitation in scanning microscopes has enabled routine *in vivo* imaging in the neocortices of rodents performing behavioral tasks.

The mammalian neocortex is often considered to have six laminar layers, demarcated by anatomical features but also populated by neurons belonging to different genetic, projection, and functional types. The cortical column is often considered a basic architecture of cortical computation, but how the six layers work in concert remains poorly understood. While two-photon excitation can easily access layer 2/3 and in a well-engineered system, down to layer 5, imaging layer 6 ($> 750\ \mu$m) in the rodent neocortex remains challenging. Given the anisotropic nature of the brain, most of the ballistic photons are scattered before reaching the intended depth, creating a background distinct in both spatial and temporal scales and drowning out the neural signals. This is largely why people have refrained from imaging these depths, even though there is evidence obtained through alternative non-imaging modalities suggesting that layer 6 neurons are responsible for modulating signal amplitudes [5], which is very distinct from the roles played by other layers.

Strategies for accessing layer 6 in the rodent neocortex have largely involved novel experimental techniques, such as single-pulse-per-pixel excitation and three-photon excitation. These techniques often require expensive bespoke femtosecond laser sources, and are challenging to set up both electronically and optically. What remains relatively under-explored are computational strategies.

A number of neural network-based denoising approaches exist in the literature, but they are generally based on Noise2Noise [3] or its derivative Noise2Self [1, 4, 2], which are intended for mean-zero shot noise removal. These are used with success on data from shallower layers, but they fall short on layer 6 data. This is most likely because the large fluorescence background as a result of tissue scattering is not mean-zero and is generally slower temporally and larger-scale spatially. The scattered background also comes with its own mean-zero shot noise component. We have optical access to the shallower layers, which is the source of the scattered background, but leveraging this information is nontrivial.

2 Overview of the project

Our goal is to reconstruct the background due to scattering, which can then be subtracted from the raw image. What is left is a movie with neuronal activities and mean-zero shot noise, which can be denoised by a more conventional Noise2Noise framework. The training of this neural network would leverage the fact that the background is larger in spatial scale and slower in temporal scale. We hope to utilize images taken from shallower layers in the neocortex–the scatterers–to aid the reconstruction of the resulting background in deeper layers. Given that cell bodies are around 10 μm in diameter and layer 6 starts around 750 μm in the mouse neocortex, one needs to sample roughly 70 to 80 planes to capture all the sources responsible for the scattering in deeper brain

layers. This is experimentally viable but time consuming. We would like to variably sample from a full z-stack and quantify how that affects the quality of reconstruction of the background.

In what follows, we briefly discuss the various aspects of this project: simulated / experimentally collected datasets, computational / experimental validation, and the timeline.

2.1 Datasets

We would like to start with simulated neuronal movies, and add various sources of noise to the simulations. We have collected movies in mice expressing CamKII-promoted GCamP8m in layer 6, which is the layer of interest, as well as data from shallower layers. The data from shallower cortical layers may be helpful as layer 5 experiences the same background but to a less extent, while layers 1-4 show very little such background.

2.2 Validation

Computational validation on simulated data is straightforward: we would compute metrics such as areas under the receiver-operation characteristic (ROC) curve by comparing extracted calcium events against the ground truths that we use to generate the simulated movies.

Experimental validation is perhaps less straightforward. So far, we have come up with two strategies. First, we would use recordings from a shallower layer that nevertheless contains significant amount of scattered background but not enough to detrimentally obfuscate the extraction of neuronal traces, and compare the traces from denoised and raw movies. Second, we would utilize techniques such as PSF engineering, and rapidly switch between a diffraction-limited PSF and a PSF prone to introduce scattering (such as a bessel beam), so one of this interleaved pair can be used as ground truth while the other the data to be tested. We have access to a deformable mirror whose switching speed is significantly faster than the calcium dynamic. Given the relative high frame rate of this system, these two movies can be considered to take place concurrently.

2.3 Timeline

We plan to create a U-net to reconstruct the background, leveraging the larger spatial (and possibly temporal) scale of the background compared to the signal. We would also like to use clean data from earlier layers to construct a prior for the clean signal. We would then include this as a regularization term in the loss function used to train the U-net for the background signal. In weeks 8 and 9, we will have finished generating simulation data, and have written code to train and test both simulated and experimentally collected movies. In week 10, we will have collected experimental validation data.

3 Goals and outlooks

The measure of success for this project is straightforward. There are two goals: (1) traces from denoised movies should closely resemble the ground truth and show no hallucination; they should also not exhibit different characteristic temporal shapes. (2) There should be significant increase in the signal-to-noise ratio (SNR) as well as signal-to-background ratio (SBR). Some biological validation could also be done. For instance, decoding the animal behavior from traces of a neural ensemble should lead to similar conclusions as decoders from raw traces (the only improvement being the enhanced detectability of features).

Not only do we expect the successful completion of this project to benefit layer 6 (and even subcortical regions) imaging with conventional experimental setups, we also think that this technique could benefit one-photon calcium imaging. One-photon imaging has many benefits, including relatively cheap illumination sources, much higher excitation efficiency, and hence easy access to large fields of view. However, the higher frequency excitation light ($\approx 1\lambda$ as opposed to $\approx 2\lambda$) leads to increased scattering and consequently low imaging depths. Our intended technique could in principle correct for scattering in one-photon movies as well, extending the imaging depths and significantly increasing the number of neurons concurrently recorded in deeper cortical layers.

References

[1] J. Batson and L. Royer. Noise2self: Blind denoising by self-supervision, 2019.

[2] J. Lecoq, M. Oliver, J. H. Siegle, N. Orlova, P. Ledochowitsch, and C. Koch. Removing independent noise in systems neuroscience data using deepinterpolation. *Nature Methods*, 18(11):1401–1408, 2021.

[3] J. Lehtinen, J. Munkberg, J. Hasselgren, S. Laine, T. Karras, M. Aittala, and T. Aila. Noise2noise: Learning image restoration without clean data, 2018.

[4] X. Li, G. Zhang, J. Wu, Y. Zhang, Z. Zhao, X. Lin, H. Qiao, H. Xie, H. Wang, L. Fang, and Q. Dai. Reinforcing neuron extraction and spike inference in calcium imaging using deep self-supervised denoising. *Nature Methods*, 18(11):1395–1400, 2021.

[5] M. Vélez-Fort, E. F. Bracey, S. Keshavarzi, C. V. Rousseau, L. Cossell, S. C. Lenzi, M. Strom, and T. W. Margrie. A circuit for integration of head- and visual-motion signals in layer 6 of mouse primary visual cortex. *Neuron*, 98(1):179–191.e6, 2026/02/19 2018.