Deep Learning for the Connectome

Seymour Knowles-Barley[1,2], Thouis R. Jones[1,2], Josh Morgan[2], Dongil Lee[2], Narayanan Kasthuri[2], Jeff W. Lichtman[2] and Hanspeter Pfister[1].
[1]Harvard University, School of Engineering and Applied Sciences, USA. [2]Harvard University, Department of Molecular and Cellular Biology, USA.

HARVARD School of Engineering and Applied Sciences
CENTER FOR BRAIN SCIENCE

Introduction

The connectome is the wiring diagram of the nervous system. Mapping this network is necessary for discovering the underlying architecture of the brain and investigating the physical underpinning of cognition, intelligence, and consciousness [1, 2, 3]. It is also an important step in understanding how connectivity patterns are altered by mental illnesses, learning disorders, and age related changes in the brain.

In our experiments, brain tissue is microtomed and imaged with an electron microscope at a resolution of 4x4x30nm per voxel. Cell membranes and sub-cellular structures such as mitochondria, synapses and vesicles are clearly visible in the images.

We have developed a fully automatic segmentation pipelines [2] to annotate this data. The requirement of a very high level of accuracy and data rates approaching 1TB a day make this a challenging computer vision task.

Deep Learning

The first and most critical stage of the our segmentation pipeline (www.rhoana.org) is to identify membrane in the 2D images produced by the microscope. We use convolutional neural networks to classify images, trained on a few hand-annotated images. GPU computing greatly reduces the training and classification time for these computationally demanding networks. In this work, we used pylearn2 [5] to apply maxout networks [6] to two connectome data sets.

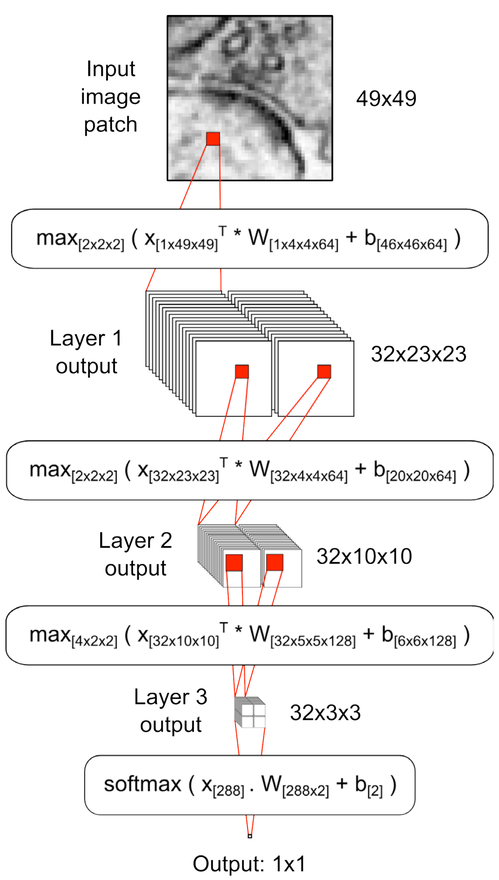

Figure 1: Deep learning maxout network structure. 49x49 pixel inputs are classified by the 4-layer deep network. Three maxout layers are used with a fully-connected softmax final layer. Networks were trained on 100k image patches with 2-class (membrane or non-membrane) labels for LGN data and 3-class (inside, membrane, outside) labels for ECS.

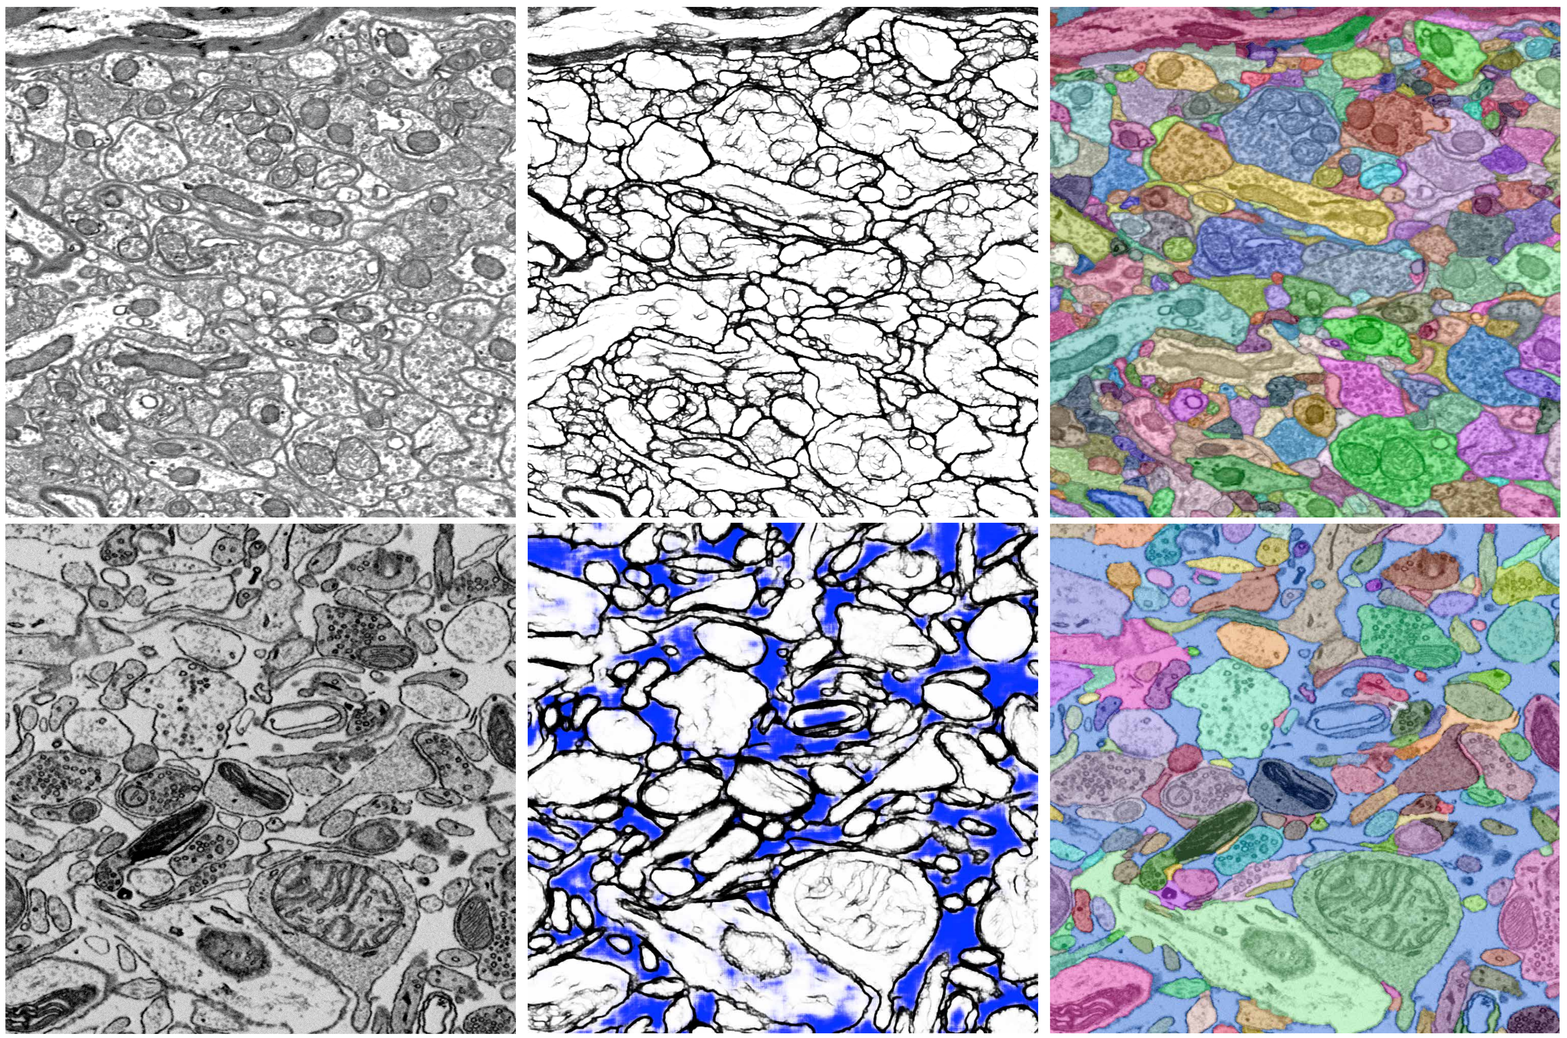

Figure 2: Left: Scanning electron microscope images of mouse lateral geniculate nucleus (LGN) (top, 8x8μm) and cortex with extra-cellular space (ECS) preserved (bottom, 4x4μm). Middle: Deep network output; membrane class shown in black and ECS class in blue. Right: Rhoana 3D segmentation pipeline output.

Results and Performance

Image classification and Rhoana segmentation results are shown in Figure 2. Pixel-wise error rates of 4.27% for 2-class LGN and 4.85% for 3-class ECS were achieved during deep network training on down-sampled, 8nm per pixel input.

Deep network classification is a computationally demanding task requiring millions of floating point operations per pixel output. We benchmark Tesla C2070, Quadro K5000 and Tesla K40c GPU performance and estimate the number of GPUs required to classify images as they are captured by a scanning electron microscope (approximately 0.85TB per day).

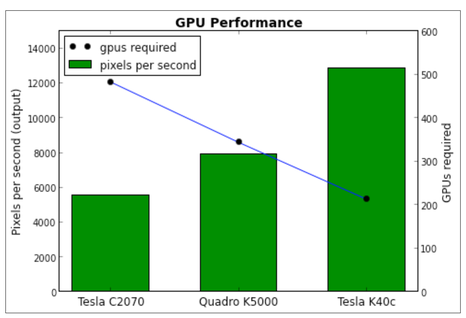

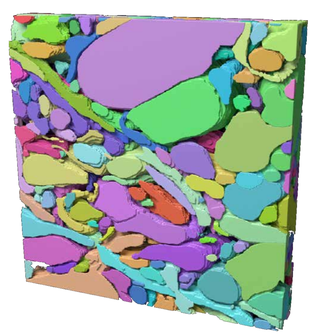

Acknowledgements
Thanks to Linda Xu, and students from Masconomet Regional High School for their careful and diligent annotations and to Daniel Berger for 3D rendering.

References

[1] Narayanan Kasthuri and Jeff W. Lichtman. Neurocartography. Neuropsychopharmacology, 35(1):342–343, 2010.
[2] Jeff W. Lichtman and Joshua R. Sanes. Ome sweet ome: what can the genome tell us about the connectome? Current Opinion in Neurobiology, 18(3):346–353, 2008.
[3] H Sebastian Seung. Reading the book of memory: Sparse sampling versus dense mapping of connectomes. Neuron, 62(1):17–29, 2009.
[4] Verema Kaynig, Amelio Vazquez-Reina, Seymour Knowles-Barley, Mike Roberts, Thouis R. Jones, Narayanan Kasthuri, Eric Miller, Jeff W. Lichtman, and Hanspeter Pfister. Large Scale Automatic Reconstruction of Neuronal Processes from Electron Microscopy Images. IEEE Transactions on Medical Imaging (under review).
[5] Ian J. Goodfellow, David Warde-Farley, Pascal Lamblin, Vincent Dumoulin, Mehdi Mirza, Razvan Pascanu, James Bergstra, Frédéric Bastien, and Yoshua Bengio. “Pylearn2: a machine learning research library”. arXiv preprint arXiv:1308.4214
[6] Ian J. Goodfellow, David Warde-Farley, Mehdi Mirza, Aaron Courville, Yoshua Bengio. Maxout Networks. ICML; JMLR WCP 28 (3): 1319-1327, 2013.

NVIDIA CUDA CENTER OF EXCELLENCE