


Decrappifying brain images with deep learning

January 8 2020, by Aaron Dubrow



An image showing the side by side versions of electron microscope captures.
Credit: Salk Institute

Textbook descriptions of brain cells make neurons look simple: a long spine-like central axon with branching dendrites. Taken individually, these might be easy to identify and map, but in an actual brain, they're more like a knotty pile of octopi, with hundreds of limbs intertwined.

This makes understanding how they behave and interact a major challenge for neuroscientists.

One way that researchers untangle our neural jumble is through microscopic imaging. By taking photographs of very thin layers of a brain and reconstructing them in three-dimensional form, it is possible to determine where the structures are and how they relate.

But this brings its own challenges. Getting high-resolution images, and capturing them quickly in order to cover a reasonable section of the brain, is a major task.

Part of the problem lies in the trade-offs and compromises that any photographer is familiar with. Open the aperture long enough to let in lots of light and any motion will cause a blur; take a quick image to avoid blur and the subject may turn out dark.

But other problems are specific to the methods used in brain reconstruction. For one, high-resolution brain imaging takes a tremendously long time. For another, in the widely-used technique called serial block face electron microscopy, a piece of tissue is cut into a block, the surface is imaged, a thin section is cut away and the block then imaged again; the process is repeated until completion. However, the electron beam that creates the [microscopic images](#) can actually cause the sample to melt, distorting the subject it is trying to capture.

Uri Manor, director of the Waitt Advanced Biophotonics Core Facility at the Salk Institute for Biological Studies in San Diego, is responsible for running numerous high powered microscopes used by researchers across the nation. He is also tasked with identifying and deploying new microscopes and developing solutions that can address problems that today's technologies struggle with.

"If someone comes with a problem and our instruments can't do it, or we can't find one that can, it's my job to develop that capability," Manor said.

Aware of the imaging issues facing neuroscientists, he decided a new approach was necessary. If he had reached the physical limits of microscopy, Manor reasoned, maybe better software and algorithms could provide a solution.

"There are sophisticated mathematical and computational approaches that have been studied for decades to remove noise without removing signal," Manor said. "That was where I started."

Working with Linjing Fang, an image analysis specialist at Salk, they cooked up a strategy to use GPUs (graphics processing units) to accelerate microscopic image processing.

They started with an image processing trick called deconvolution that had been developed in part by John Sedat, one of Manor's scientific heroes and a mentor at Salk. The approach was used by astronomers who wanted to resolve images of stars and planets with greater resolution than they could achieve directly from telescopes.

"If you know the optical properties of your system, then you can deblur your images and get twice the resolution of the original," he explained.

They believed that deep learning—a form of machine-learning that uses multiple layers of analysis to progressively extract higher level features from raw input—could be very useful for increasing the resolution of microscope images, a process called super-resolution.

MRIs, [satellite imagery](#), and photographs had all served as [test cases](#) to develop deep learning-based, super-resolution approaches, but

remarkably little had been done in microscopy. Perhaps, Manor thought, the same could be done with microscopy.

The first step in training a deep learning system involves finding a large corpus of data. For this, Manor teamed up with Kristen Harris, a neuroscience professor at The University of Texas at Austin and one of the leading experts in brain microscopy.

"Her protocols are used around the world. She was doing open science before it was cool," Manor said. "She gets incredibly detailed images and has been collaborating with Salk for a number of years."

Harris offered Manor as much data as he needed for training. Then, using the Maverick supercomputer at the Texas Advanced Computing Center (TACC) and several days of continuous computation, he created low-resolution analogs of the high-resolution microscope images and trained a deep learning network on those image pairs.

"TACC has been incredibly helpful," Manor said. "They provided us with hardware to do training before our hair fell out and provided us with computational expertise and even helped run computational experiments to fine-tune our process."

Unfortunately, Manor's first attempts to create super-resolution versions of low-resolution images were unsuccessful. "When we tried to test the system on real world low resolution data that was much noisier than our low resolution training data, the network didn't do so well."

Manor had another stroke of luck when Jeremy Howard, founder of fast.ai, and Fred Monroe, from the Wicklow AI Medical Research Initiative (WAMRI.ai), came to Salk looking for research problems that could benefit from deep learning.

"They were excited by what we doing. It was a perfect application for their deep learning methods and their desire to help bring deep learning to new domains," Manor recalled. "We started using some of their tricks that they had established, including crappification."

At the time of their meeting, Manor and Fang had been computationally decreasing the resolution of their images for training pairs, but they were still not crappy enough. They were also using a type of deep learning architecture called generative adversarial networks (GANs).

"They suggested adding more noise computationally," he recalled. "'Throw in some blur, and different kinds of noise, to make images really crappy.' They had built a library of crappifications and we crappified our images until it looked much more like, or even worse than, what it looks like when you acquire a low resolution image in the world. They also helped us switch away from GANs to U-Net architectures, which are much easier to train and better at removing noise."

Manor retrained his AI system using the new image pairs and deep learning architecture and found that it could create [high-resolution images](#) that were very similar to the ones that had been originally created with greater magnification. Moreover, trained experts were able to find brain cell features in decrappified versions of the low-res samples that couldn't be detected in the originals.

Finally, they put their system to the real test: applying the method to images created in other labs with different microscopes and preparations.

"Usually in deep learning, you have to retrain and fine tune the model for different data sets," Manor said. "But we were delighted that our system worked so well for a wide range of sample and image sets."

The success meant that samples could be imaged without risking damage, and that they could be obtained at least 16 times as fast as traditionally done.

"To image the entire brain at full resolution could take over a hundred years," Manor explained. "With a 16 times increase in throughput, it perhaps becomes 10 years, which is much more practical."

The team published their results in [Biorxiv](#), presented them at the F8 Facebook Developer Conference and the 2nd NSF NeuroNex 3DEM Workshop, and made the code available through GitHub.

"Not only does this approach work. But our training model can be used right away," Manor said. "It's extremely fast and easy. And anyone who wants to use this tool will soon be able to log into 3DEM.org [a web-based research platform focused on developing and disseminating new technologies for enhanced resolution 3-dimensional electron microscopy, supported by the National Science Foundation] and run their data through it."

"Uri really fosters this idea of image improvement through [deep learning](#)," Harris said. "Ultimately, we hope we will not have any crappy images. But right now, many of the images have this problem, so there's going to be places where you want to fill in the holes based on what's present in the adjacent sections."

Manor hopes to develop software that can do reconstruction on the fly, so researchers can see super resolution images right away, rather than in post-processing. He also sees the potential for improving the performance of the millions of microscopes already at labs around the world and for building a brand new microscope from the ground up that takes advantage of AI capabilities.

"Less expensive, higher resolution, faster—there are lots of areas that we can improve upon."

With a proof of concept in place, Manor and his team have developed a tool that will enable advances throughout neuroscience. But without fortuitous collaborations with Kristen Harris, Howard and Monroe and TACC, it may never have come to fruition.

"It's a beautiful example of how to really make advances in science. You need to have experts open to working together with people from wherever in the world they may be to make something happen," Manor said. "I just feel so very lucky to have been in a position where I could interface with all of these world-class teammates."

More information: www.biorxiv.org/content/biorxiv/2020/06/07/740548.full.pdf

Provided by University of Texas at Austin

Citation: Decrappifying brain images with deep learning (2020, January 8) retrieved 23 April 2024 from <https://techxplore.com/news/2020-01-decrappifying-brain-images-deep.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.