



ILLUSTRATION BY THE PROJECT TWINS

AI SPOTS CELL STRUCTURES THAT HUMANS CAN'T

Models can predict the location of cell structures from light-microscopy images alone, without using harmful fluorescence labels. **By Amber Dance**

Susanne Rafelski and her colleagues had a deceptively simple goal. “We wanted to be able to label many different structures in the cell, but do live imaging,” says the quantitative cell biologist and deputy director of the Allen Institute for Cell Science in Seattle, Washington. “And we wanted to do it in 3D.”

That kind of goal normally relies on fluorescence microscopy – problematic in this case because, with only a handful of colours to use, the scientists would run out of labels well before they ran out of structures. Also problematic is that these reagents are pricey and laborious to use. Moreover, the stains are harmful to live cells, as is the light used to stimulate them, meaning that the very act of imaging cells can damage them. “Fluorescence is expensive, in many different versions of the word ‘expensive,’” says Forrest Collman, a microscopist at the Allen Institute for Brain

Science, also in Seattle. When Collman and his colleagues tried to make a 3D time-lapse movie using three different colours, the results were “horrific”, Collman recalls. “The cells all just die in front of you.”

Imaging cells using transmitted white light (bright-field microscopy) doesn’t rely on labelling, so avoids some of the problems of fluorescence microscopy. But the reduced contrast can make most cell structures impossible to spot. What Rafelski’s team needed was a way to combine the advantages of both techniques. Could artificial intelligence (AI) be used on bright-field images to predict how the corresponding fluorescence labels would look – a type of ‘virtual staining’? In 2017, Rafelski’s then-colleague, machine-learning scientist Gregory Johnson, proposed just such a solution: he would use a form of AI called deep learning to identify hard-to-spot structures in bright-field images of unlabelled cells.

“No way,” said Rafelski, as she headed off for a few months’ leave. When she returned to work, Johnson told her he’d done it. “It blew my mind that it was possible,” Rafelski recalls. Using a deep-learning algorithm on unlabelled cells, the Allen team created a 3D film showing DNA and substructures in the nucleus, plus cell membranes and mitochondria¹.

“These models are ‘seeing’ things that humans don’t,” says Jason Swedlow, a quantitative cell biologist at the University of Dundee, UK. Our eyes, he says, just aren’t adapted to pick out subtle, greyscale patterns such as those in optical microscopy – that’s not how we evolved. “Your eyes are supposed to see lions and trees and things like that.”

Over the past few years, scientists working on AI have designed several systems that can pick out those patterns. Each model is trained using pairs of images of the same cells, one bright-field and one fluorescently labelled.

But the models differ in the details: some are meant for 2D images, some for 3D; some aim to approximate cellular structures whereas others create pictures that could be mistaken for true photomicrographs.

“This represents a huge advance in what we are able to achieve,” says Mark Scott, microscopy facility manager at the Translational Research Institute Australia in Brisbane. What’s needed now is for biologists to collaborate with the AI coders, testing and improving the technology for real-world use.

Fast-growing field

Steven Finkbeiner, a neuroscientist at the University of California, San Francisco, and the Gladstone Institutes, also in San Francisco, uses robotic microscopy to track cells for up to a year. By the early 2010s, his group was accumulating terabytes of data per day. That caught the attention of researchers at Google, who asked how they might help. Finkbeiner suggested using deep learning to find the cellular features he couldn’t see.

Deep learning uses computer nodes layered in a similar way to neurons in the human brain. At first, the connections between nodes in this neural network are weighted randomly, so the computer is only guessing. But with training, the computer adjusts the weights, or parameters, until it starts to get it right.

Finkbeiner’s team trained its system to identify neurons in 2D images, then pick out the nucleus and determine whether a given cell is alive or not². “The main point was to show scientists that there is probably a lot more information in image data than they realize,” says Finkbeiner. The team called its technique *in silico* labelling.

The approach couldn’t identify motor neurons, however – perhaps because there wasn’t anything in the unlabelled cells to indicate their specialization. These predictions will only work if there’s some visible cue that the AI can use, Collman says. Membranes, for example, have a different refractive index to their surroundings, producing contrast.

Collman, Johnson and their colleagues at the Allen Institute used a different neural network to solve Rafelski’s problem, building on a system called U-Net that was developed for biological images. Unlike Finkbeiner’s approach, the Allen model works with 3D micrographs, and some researchers at the institute now use it routinely – for example, to identify nuclear markers in studies of chromatin organization.

At the University of Illinois at Urbana-Champaign, physicist Gabriel Popescu is using deep learning to answer, among other things, one of the most fundamental microscopy questions: is a cell alive or dead? That’s harder than it sounds because tests for life, paradoxically, require toxic chemicals. “It’s like taking the pulse of the patient with a knife,” he says.

Popescu and his colleagues call their

approach PICS: phase imaging with computational specificity. Popescu uses it in live cells to identify the nucleus and cytoplasm, then calculates their masses over days at a time³. These signatures accurately indicate cell growth and viability, he says.

PICS encompasses software based on U-Net and microscope hardware, so instead of obtaining images and training a machine to process them later, it all happens seamlessly. Once a user snaps a white-light image, it takes just 65 milliseconds for the model to deliver the predicted fluorescence counterpart.

Other groups use different kinds of machine learning. For instance, a team at the Catholic University of America in Washington DC used a type of neural network called a GAN to identify nuclei in images from phase-contrast optical microscopy⁴. A GAN, or generative adversarial network, sets up two opposing models: the ‘generator’ predicts the fluorescence images, and the ‘discriminator’ guesses whether they’re real or fake. When the discriminator is fooled about half the time, the generator must be making plausible predictions, says Lin-Ching Chang, an engineer on the project. “Even humans cannot tell the generated examples are fake.”

Drug discovery

Fluorescence predictions are also taking hold in the drug industry. At AstraZeneca in Gothenburg, Sweden, pharmacologist Alan Sabirsh studies fat cells for their roles in disease and drug metabolism. Sabirsh and AstraZeneca teamed up with the Swedish National Center for applied Artificial Intelligence to run the Adipocyte Cell Imaging Challenge, asking competitors to identify the nucleus, cytoplasm and lipid droplets in unlabelled micrographs. Its US\$5,000 prize went to a team led by Ankit Gupta and Håkan Wieslander, two PhD students at Uppsala University in Sweden, who work on image processing.

Like Chang and her colleagues, the team used a GAN to identify lipid droplets. But to get at the nuclei, they used a different technique, called LUPi – learning using privileged information, which gives the machine extra help as it learns. In this case, the team used a further image-processing technique to identify the nuclei in the standard training image pairs. Once the model was trained, however, it could predict nuclei on the basis of light-microscopy images alone⁵.

The resulting images aren’t perfect: Gupta says real fluorescence staining provides more realistic texturing in the nucleus and cytoplasm than the model can. It’s good enough for Sabirsh, however. He has already started using the code in robotic-microscopy experiments with the aim of developing therapeutics.

With several proof-of-principle projects complete, the technique has moved beyond the first baby steps, says Swedlow, and the

wider community is beginning to put it through its paces. “I think we are learning to walk, and what it means to walk,” he says.

For example, when is it beneficial to make predictions on the basis of white-light images, and when should it be avoided? Trying to determine segmentation of cellular compartments and structures is probably a good application, because any errors won’t significantly affect downstream results, says Anne Carpenter, senior director of the Imaging

“Even humans cannot tell the generated examples are fake.”

Platform at the Broad Institute of MIT and Harvard in Cambridge, Massachusetts. She’s more circumspect about predicting experimental outcomes, however, because the machine might rely on one structure that predicts another only under control conditions. “Often, in biology, it’s the exceptions to the rule that are what we’re looking for,” Carpenter says.

For now, at least, scientists would do well to confirm a model’s key predictions using standard fluorescence staining, says Popescu. And it’s a good idea to seek expert collaborators, adds Laura Boucheron, an electrical engineer at New Mexico State University in Las Cruces. “There’s a lot of very significant computer know-how required to even get these up and running.”

Some models use just a handful of images for training, but Boucheron cautions that larger data sets are preferable. Hundreds, or better yet thousands, might be required, says Yvan Saeys, a computational biologist at the VIB Center for Inflammation Research at Ghent University in Belgium. And if you want the model to work with multiple cell types or different microscope set-ups, be sure to include that variety in the training set, he adds.

Large-volume training might require weeks of time on supercomputers with multiple graphical processing units, warns Boucheron. But once that’s done, the prediction model could run off a laptop, or even a mobile phone.

For many researchers, that one-time investment is worth it, if it means never staining for this or that feature again. “If you could collect pictures of unlabelled cells and you already had trained algorithms,” says Finkbeiner. “You get all that information, basically, for free.”

Amber Dance is a science journalist in Los Angeles, California.

1. Ounkomol, C., Seshamani, S., Maleckar, M. M., Collman, F. & Johnson, G. R. *Nature Meth.* **15**, 917–920 (2018).
2. Christiansen, E. M. et al. *Cell* **173**, 792–803 (2018).
3. Kandel, M. E. et al. *Nature Commun.* **11**, 6256 (2020).
4. Nguyen, T. C. et al. *J. Biomed. Opt.* **25**, 096009 (2020).
5. Wieslander, H., Gupta, A., Bergman, E., Hallström, E. & Harrison, P. J. Preprint at bioRxiv <https://doi.org/10.1101/2021.01.18.427121> (2021).