Yair Rivenson and Aydogan Ozcan

Toward a Thinking Microscope

Convolutional neural networks and deep learning can boost the capabilities of standard optical microscopes to levels comparable to those of some higher-end imaging systems.

Illustration by A. Ozcan

Deep learning, particularly using convolutional neural networks (CNNs), is transforming a range of disciplines and eclipsing the state of the art achieved by earlier machine-learning techniques. In machine vision, for example, the deep-learning revolution has driven new capabilities in autonomous vehicles, fault analysis, security applications, entertainment and the Industrial Internet of Things. Deep-learning-enabled breakthroughs in voice recognition and speech translation are transforming how we communicate with each other and with our devices. And supervised deep-learning approaches—in which a system learns to classify or otherwise interpret information by analyzing "training sets" of labeled data—have found particular use in biomedicine and medical imaging: disease diagnosis through classification of histological images; determining tumor margins in cancer cases; cell classification and counting; screening patients for certain eye diseases using optical coherence tomography scans; and a host of other areas.

Beyond these mainstream applications related to classifying and interpreting microscope images, supervised deep learning has created new opportunities to reconstruct and improve the quality of the images themselves. This article looks at some recent efforts to apply state-of-the-art deep-learning approaches to optical microscopy, microscopic image reconstruction, and image transformation in general. Deep-learning-enabled reconstruction and transformation of optical-microscopy images acquired with certain imaging systems can significantly boost their resolution, field of view and depth of field, and can correct for various sources of aberrations. As a result, the raw images from relatively low-end microscopy systems can, with the aid of deep learning, be transformed to match the images expected from a much higher-end system.

Training CNNs with image data

The deep-learning revolution is the result of a "perfect storm" that has revolutionized computation generally: access to the enormous amounts of data generated by digital sensors, imagers and consumer devices, coupled with the availability of powerful yet cost-effective hardware and software tools, which make large-scale statistical learning based on those data volumes a tractable challenge. Using deep learning for reconstruction or transformation of images finds particularly fertile ground in optical microscopy, since, relative to photography or macro-scale imaging in general, optical microscopy provides a highly controlled, repeatable platform. That, in turn, can allow the creation of nanoscopically aligned and labeled image pairs (for example, matched pairs showing the same sample viewed in a low-resolution or aberrated microscope and a higher-resolution, diffraction-limited one) that can be rolled into sets of images for training the deep-learning system.

Following the collection of the matched pairs of training data, which typically contain thousands of image patches, the data are fed into a model that learns the statistical image transformation from an input distribution (for example, that of the low-resolution microscope) to the desired or enhanced output distribution (represented by the images from the higher-end instrument). For deep neural networks, the model architecture commonly takes the form of multiple, hierarchical cascaded convolutional layers. Each convolutional layer contains multiple convolution kernels and bias terms (adjustable parameters that can be trained) that act as filters, operating on units of the convolutional layer known as feature maps, or channels.

The resulting weighted sum from the filtered feature map next passes through a *nonlinear activation function*, the output of which is sent to the next convolutional layer. This nonlinear function, which can take a number of mathematical forms, is at least intuitively analogous to the activation functions at the synapses of biological neurons, where the decision triggering a neuron firing event is made based on whether the "synaptic weight" exceeds a particular threshold. The nonlinear nature of the activation function helps the network to learn complex transformations and generalize its inference to a broader class of functions or tasks. The filtered output passes from one layer to another, with each layer learning a higher level of abstraction.

Generally, increasing the number convolutional filters should allow for more complex network models and, in principle, for learning more complicated tasks or relationships with larger datasets. The price, however, is an increase in the training and inference times—as well as the possibility of overfitting the model to a particular set of training data.

Trained and ready

Training a CNN based on labeled data is fundamentally an optimization problem: the network attempts to optimize its output with respect to "gold standard" target labels, given a user-defined cost (or loss) function.

1047-6938/18/07-08/34/8-$15.00 © OSA

Supervised deep learning has created new opportunities to reconstruct and improve the quality of microscopy images.

For example, one can choose to minimize the energy difference (such as the mean squared error) between the network's output images and their corresponding gold-standard labels (for instance, the image data from a higher-end microscope that the neural network is trained to mimic).

The infographic below shows a typical training procedure. For each pair of images, the amount of error is calculated through the predefined cost function; this error is propagated back within the network to update all the weights and biases in the direction that minimizes the error or the cost. Through this iterative process, known as error back-propagation in deep learning, the model adjusts its coefficients to satisfy the given target labels and minimize cost function constraints.

The process iterates on the available training image data to fine-tune the model's inference success. When the entire training dataset has been used during these iterations, it constitutes an "epoch"; typically tens to thousands of epochs are required to complete the training phase, depending on the complexity of the network and the task that it is charged with. After this phase, the network is ready to be blindly evaluated with a new "test set" of image data that were not part of the original set of training images.

An exciting aspect of deep-learning-enhanced optical microscopy and image transformations is that, once the network is trained, the inference is non-iterative and very fast to compute, without the need for parameter search or optimization. In this sense, compared with other solutions to inverse problems in optics—for example, deconvolution, convex optimization or compressive sampling or sensing techniques—a trained deep neural network, once optimized, has a significant edge in computation speed and inference time, even with modest computers and processors. That edge is getting even greater with the emergence of new processor architectures that are specifically optimized for deep learning, and could ultimately let neural networks perform their inference tasks in

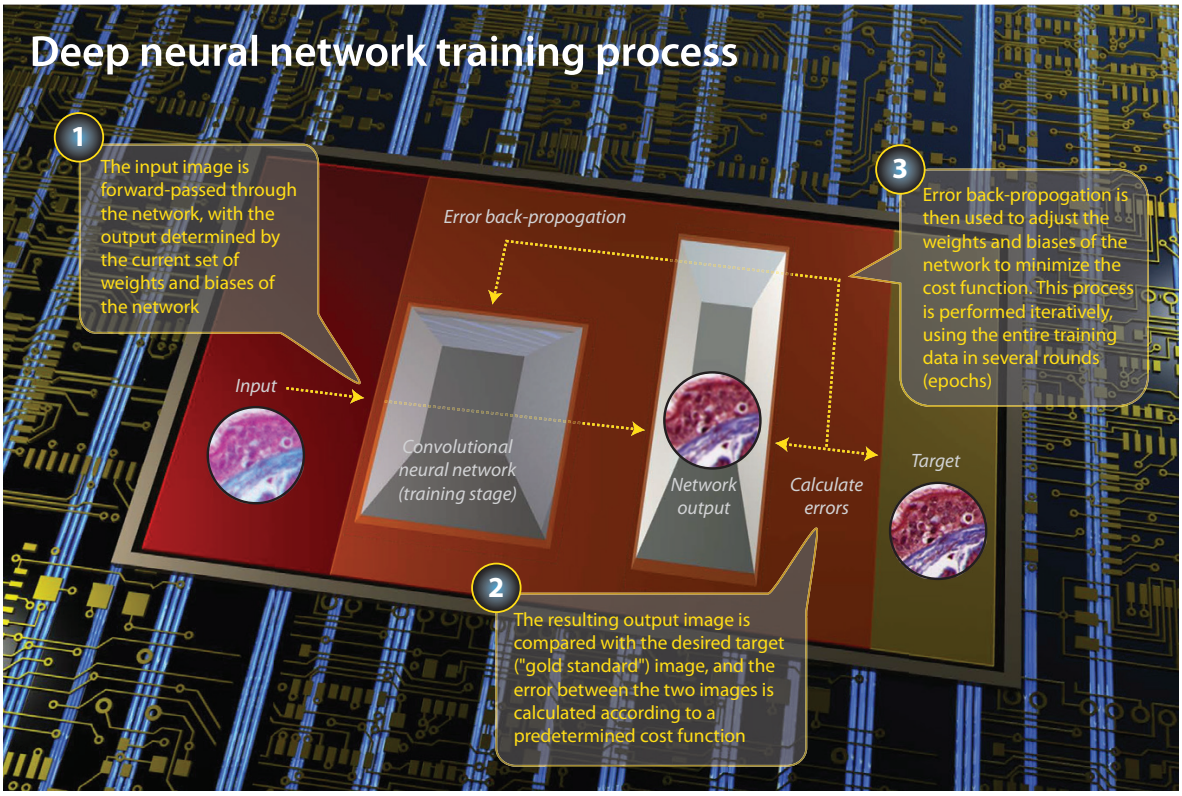

Illustration by Phil Saunders

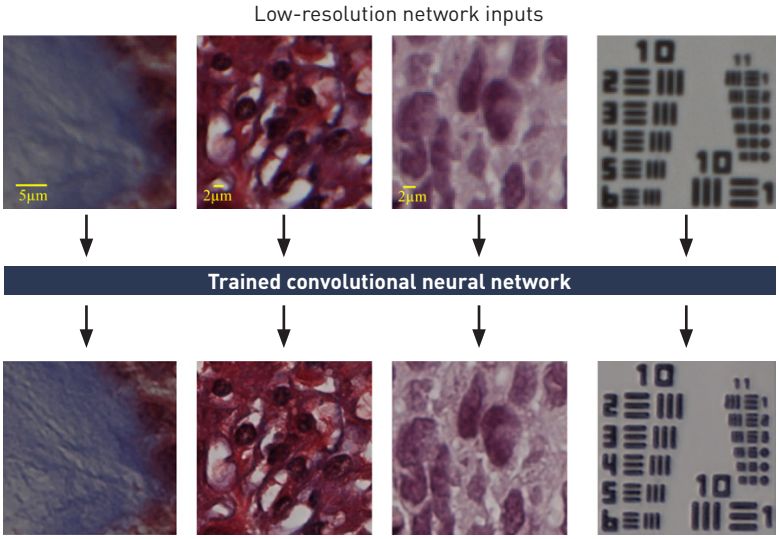

Image enhancement

A deep neural network trained only with lung tissue samples to sharpen their images (left column) also improved low-resolution images of kidney and breast tissue sections (middle columns) and a resolution test target (right column) that were never seen by the network during its training.

real time, even with mobile phones and other low-end consumer devices.

Training the network can take from a few hours to more than a day, depending on the size of the training data, the available hardware, and the complexity of the model, among other things. Once the model is trained, however, it remains fixed. Furthermore, in a process called transfer learning, a trained neural network can even be used to "warm start" new models, when new data become available or new tasks are required.

Improving bright-field microscopy

While deep learning has been used for a number of years to classify and annotate microscopy images, one of the first applications of deep learning to enhance optical microscopy images was recently demonstrated using bright-field microscopy. The training set for the effort began with histochemically stained lung tissue sections, each one imaged twice: once using a bright-field microscope with a 40×/0.95NA objective lens, to obtain lower-resolution (LR) images of specimen; and once with a 100×/1.4NA oil-immersion objective lens, used to obtain the corresponding high-resolution (HR) labels or gold-standard images. A deep neural network architecture was then designed to transform LR images (used as input) into enhanced images matching the HR labels.

Essentially, the network's goal is to predict the pixel values of the HR image, given the LR input. Therefore, an important step before training is to precisely align, or register, the LR and HR training images with respect to each other; this enforces the deep neural network to solely learn the LR-to-HR transformation, rather than some arbitrary transformation between the input and output images related, for example, to misalignment. Following the accurate alignment of the images, the network model can be trained with the matched LR and HR image pairs.

One key advantage of deep learning over other image enhancement or deconvolution methods is that no *a priori* information about the image formation process is required. That is, modelling of the point-spread function, spatial and spectral aberrations, illumination properties or other physical parameters of the imaging system or the object, and their impact on the acquired image, do not need to be known or estimated. Instead, the neural network uses training image data to inherently learn these details in its multidimensional solution space.

After the training step, the network was blindly tested on Masson's Trichrome–stained lung tissue sections taken from a different patient. The network output, in response to LR input images, super-resolved the blurry and distorted features in the input images, providing images similar to those acquired with a 100×/1.4NA objective. The trained network was also able to boost image quality for new types of samples that were not part of the training data. For example, the same model, trained with lung tissue sections, was tested on kidney tissue images stained with Masson's Trichrome, and was able to enhance the resolution of the imaged specimen. Furthermore, the same deep network could super-resolve different tissue types that used a different stain (for instance, breast tissue sections labeled with Haemotoxylin and Eosin). The network even achieved enhanced resolution and data-driven frequency extrapolation when blindly tested

One key advantage of deep learning over other image enhancement methods is that no *a priori* information about the image formation process is required.

on a standard resolution test target—a clear indication of the generalizability of the deep-learning approach.

Another interesting feature of this bright-field microscopy network is that it extends the inferred image's depth of field. Since the input images are acquired using a lower-NA objective lens, the network learns to enhance all the spatial features that appear in focus in a low-NA image, resulting in an output image with high resolution over an extended depth of field. Because the input images have a larger field of view compared with the higher-NA objective lens that the network is trained for, the network output images also exhibit increased field of view, further enhancing the imaging throughput of bright-field microscopy through deep learning.

Laptops and smartphones

Of particular importance is that, even using an ordinary laptop computer equipped with a graphics processing unit, the network output image can be calculated in less than a second, without any iterations or parameter tuning. That suggests the potential of real-time performance using more advanced computational resources and parallel computing. And, in another vein, recent work suggests that deep learning could significantly improve the imaging performance of mobile-phone-based microscopes.

The creation of cost-effective and portable microscopic imaging systems based on mobile devices such as smartphones has advanced significantly in recent years, with potentially profound effects for global health and point-of-care diagnostics. Yet mobile microscopy devices still fall short of the quality of laboratory-grade microscopes used in clinical applications, in part because of design constraints imposed by the compactness, cost-effectiveness and extremely large volume manufacturing of mobile phones.

Recent work suggested the potential of deep learning to bridge this performance gap. For this goal, as with the bright-field microscopy example, thousands of aligned image patches obtained by smartphone-based and benchtop microscopes were used to train the network. The resulting trained system was able to substantially improve the quality of subsequently acquired smartphone images, correcting various aberrations, blocking artifacts, and increasing signal-to-noise ratio and spatial resolution. Once again, the results required no numerical or analytical modeling of the spatially and spectrally varying aberrations of the mobile microscope to set-up an inverse problem based on a forward model. (In fact, such a forward model is rather complicated to establish for a mobile phone-based microscope, as the repeatability of a cost-effective handheld platform is not high.) Training of a deep neural network with image data thus provides an elegant solution for statistically learning the transformation between LR or aberrated input images and the HR labels, without any physical modeling of the image formation process.

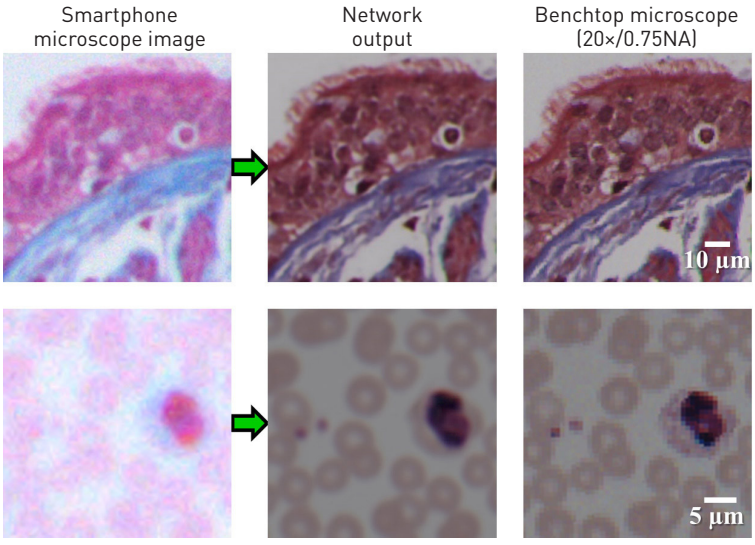

Improving smartphone-based microscopy

A trained convolutional neural network enhances smartphone-based microscope images of a lung tissue section (top row) and a blood smear (bottom row).

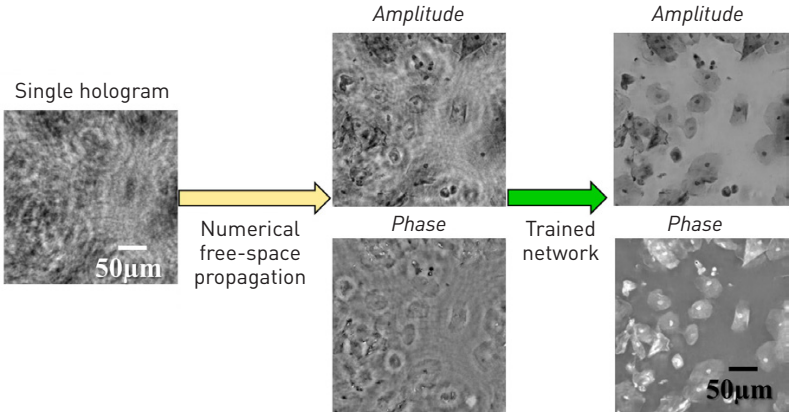

Holographic image reconstruction

Because phase information is missing in a single hologram of a Pap smear (left column), the object's complex-valued image is distorted by the twin-image and self-interference artifacts, following the free-space back propagation without phase (middle column). The deep network is trained to rapidly reconstruct the true object distribution, in both the phase and amplitude channels (right column), non-iteratively and using a single hologram of the sample.

The same work revealed that the network could work with significantly compressed input images. Lossy compressed JPEG images with a file size more than 20 times smaller than their lossless compressed counterparts, captured by the same smartphone microscope, were used as inputs to the same neural network, with very similar inference results, matching the HR labels acquired using a benchtop bright-field microscope. This ability to handle lossy compressed images could prove especially important for point-of-care applications in resource-scarce settings that have limited data transmission bandwidth and storage capacity.

Opportunities in holography

In addition to bright-field microscopy, deep learning has also been applied to improve other optical microscopy modalities, including fluorescence microscopy and holographic microscopy. For fluorescence microscopy, several recent approaches have focused on accelerating the image acquisition phase of localization-based super-resolution microscopy, using numerical models of the imaging system. Another recent approach has considered learning of the statistical image transformation between low-resoution and high-resolution datasets, without any prior assumptions on the imaging system. Similar to fluorescence microscopy, holography exemplifies another unique opportunity, where deep-learning-based methods open up data-driven alternatives to decades-old analytical or iterative, physics-driven microscopic image reconstruction techniques.

Holographic microscopy can indirectly detect both the amplitude and phase of the light field from an object. While coherent holographic imaging systems offer some unique advantages for label-free sample analysis, they generally suffer from the "missing phase" problem: missing phase information at the detector plane means that the numerically formed object image, unless the phase is recovered, will be plagued by artifacts such as self-interference noise and twin images. The latter is especially strong for in-line holographic microscopy, where the object wave and the reference wave co-propagate in the same direction.

A number of phase recovery and holographic image reconstruction approaches have been developed over the years to solve the missing-phase problem. Purely algorithmic approaches require prior information on the object function, such as its spatial support or sparsity representation. Other approaches involve hardware modifications of the holographic imaging setup to facilitate measurement diversity, and the use of additional measurements (via, for example, multiple angles of illumination, multiple sample-to-sensor distances, phase-shifting or other approaches) as physical constraints for phase recovery and image reconstruction.

As an alternative, data-driven methods based on deep learning can perform holographic image reconstruction from a single hologram, providing significant savings in both hologram acquisition and reconstruction times. One such deep-learning model was trained using numerically back-propagated in-line hologram intensities (without phase information) as complex-valued inputs to the network. The target images—that is, the labels in the supervised deep-learning scheme—were reconstructed by an iterative multi-height phase retrieval method that used eight different holographic measurements of the same object, taken at different sample-to-sensor distances, providing physical measurement diversity for accurate phase retrieval.

After its training, this phase recovery and holographic image reconstruction network was blindly

The marriage of optical imaging with *unsupervised* deep learning could bring us closer to a thinking microscope.

tested by imaging various samples, including breast tissue sections, blood and Papanicolaou (Pap) smears, using a single hologram measurement in each case. The images reconstructed by deep learning very well matched gold-standard images obtained using eight holograms of the same sample, processed with the iterative multi-height phase retrieval algorithm. The phase recovery neural network successfully learned to eliminate self-interference and twin-image-related artifacts that are normally superimposed onto the phase and amplitude channels of the object's image.

Despite all of its training data and success in holographic reconstructions, the network did not learn the actual physics of wave propagation or the hologram formation process, as it was not trained for it. In fact, some of the out-of-focus objects (such as dust particles) that lay outside of the sample plane were also cleaned/removed by the neural network, although they are physical objects that should appear in a physics-driven hologram reconstruction method. This means that, instead of providing a solution that is compatible with the wave equation, the network only learned the image transformation that it was statistically trained for.

Another benefit of this deep-learning-based holographic image reconstruction approach is its rapid reconstruction time—at least four times faster than iterative reconstruction approaches. More recent work suggests that using a similar deep-learning framework to perform both autofocusing and phase recovery tasks in a single neural network can provide significantly larger depth of field in holographic imaging, while also dramatically improving the algorithm time-complexity of holographic image reconstruction in general.

A paradigm shift

As the examples above suggest, deep learning has the potential to change the nature of optical microscopy and image reconstruction methods, by enabling new transformations among different modes and modalities of microscopic imaging—all driven by image data. We believe that deep learning will become an essential component of modern optical microscopy, fundamentally changing both its hardware and image reconstruction methods in a holistic manner.

We also emphasize that the examples presented here have involved supervised deep learning, in which known data labels are used as feedback to train a statistical model. Some experts in deep learning believe that *unsupervised* deep learning represents a very important aspect of the discipline's future. Unsupervised learning does not rely on label-based feedback; instead, the network attempts to infer the internal structure of the input data, often through some policy of rewards. The emergence of powerful systems that utilize unsupervised deep learning approaches, still in their infancy, could, we believe, ultimately lead to a sort of learning and thinking microscope—with unique capabilities that will enable applications not possible with today's optical microscopy technologies. OPN

Yair Rivenson and Aydogan Ozcan (ozcan@ucla.edu) are with the Electrical and Computer Engineering Department, Bioengineering Department, and California NanoSystems Institute, University of California, Los Angeles, Calif., USA.

References and Resources

- Y. LeCun et al. "Deep learning," Nature **521**, 436 (2015).
- G. Litjens et al. "A survey on deep learning in medical image analysis," Med. Image Anal. **42**, 60 (2017).
- Y. Rivenson et al. "Deep learning microscopy," Optica **4**, 1437 (2017).
- A. Sinha et al. "Lensless computational imaging through deep learning," Optica **4**, 1117 (2017).
- N. Boyd et al. "DeepLoco: Fast 3D localization microscopy using neural networks," bioRxiv 267096, doi: 10.1101/267096 (2018).
- D.S. Kermany et al. "Identifying medical diagnoses and treatable diseases by image-based deep learning," Cell **172**, 1122 (2018).
- E. Nehme et al. "Deep-STORM: Super-resolution single-molecule microscopy by deep learning," Optica **5**, 458 (2018).
- Y. Rivenson et al. "Deep learning enhanced mobile-phone microscopy," ACS Photon., doi: 10.1021/acsphotonics.8b00146 (2018).
- Y. Rivenson et al. "Phase recovery and holographic image reconstruction using deep learning in neural networks," Light Sci. Appl. **7**, 17141 (2018).
- Y. Rivenson et al. "Deep learning–based virtual histology staining using auto-fluorescence of label-free tissue," arXiv:1803.11293 (2018).
- H. Wang et al. "Deep learning achieves super-resolution in fluorescence microscopy," bioRxiv 309641, doi: 10.1101/309641 (2018).
- Wu et al. "Extended depth-of-field in holographic image reconstruction using deep learning-based auto-focusing and phase-recovery," Optica **5**, 704 (2018).