

# Applications, Promises, and Pitfalls of Deep Learning for Fluorescence Image Reconstruction

Chinmay Belthangady<sup>1</sup>, & Loic A. Royer<sup>1\*</sup>

<sup>1</sup>Chan Zuckerberg Biohub, San Francisco, USA

**Deep Learning is a recent and important addition to the computational toolbox available for image reconstruction in fluorescence microscopy. We review state-of-the-art applications such as image restoration and super-resolution, and discuss how the latest Deep Learning research can be applied to other image reconstruction tasks such as structured illumination, spectral deconvolution, and sample stabilisation. Despite its successes, Deep Learning also poses significant challenges, has often misunderstood capabilities, and overlooked limits. We will address key questions, such as: What are the challenges in obtaining training data? Can we discover structures not present in the training data? And, what is the danger of inferring unsubstantiated image details?**

**Keywords:** microscopy, fluorescence, machine learning, deep learning, inverse problems, image reconstruction, image restoration, super-resolution, deconvolution, spectral unmixing.

Fluorescence microscopy is an indispensable tool in the Biologist's arsenal that has made possible the systematic spatio-temporal dissection of life's molecular machines<sup>1</sup>. In all microscopy modalities, the resolution and quality of the images obtained is fundamentally limited by the optics, the photo-chemistry of molecular probes, and sensor technology. Recent developments such as super-resolution and structured illumination microscopy (SIM) can circumvent these limits by means of clever experimental strategies followed by computation. The final high quality images must be *reconstructed* from raw acquisitions. In general, microscopy is becoming increasingly dependent on computation as an integral part of the imaging process. Any significant improvements in the algorithms used for reconstruction would not only improve image quality but also open the door to novel imaging modalities.

*Deep Learning*<sup>2</sup> and its stunning successes at solving hard and previously unsolved computational problems (see **Box 1**) explains why, in recent years, machine learning and in particular deep *Convolutional Neural Networks* (CNN) have impacted almost every discipline of science and engineering from Astronomy<sup>3</sup> and Biology<sup>4</sup> to High-Energy Physics<sup>5</sup>. CNNs are particularly relevant in fields that produce or interpret images. For example, medical imaging picked up the trend early on and successfully applied deep CNNs to image

processing of computed tomography and magnetic resonance images<sup>6,7</sup>. Fluorescence microscopy too, has recently benefited from these advances, with deep CNNs applied to image restoration<sup>8</sup>, deconvolution<sup>9,10</sup>, super-resolution<sup>11-13</sup>, translation between label-free and fluorescence images<sup>14,15</sup>, and 'virtual' staining<sup>16</sup>.

## Deep Learning for image reconstruction.

A common denominator between applications mentioned above is the formulation of image reconstruction as a transformation between acquired and reconstructed images. Classical approaches to image reconstruction have to be handcrafted from first principles (see **Box 2**). For example, in deconvolution, this requires precise understanding of the optics and well characterised noise statistics. This has led to the design of popular algorithms such as Richardson-Lucy deconvolution which require knowledge of the point spread function of the microscope and assumes Poisson noise statistics<sup>21</sup>. However, such handcrafted algorithms are limited by the accuracy of their assumptions. They fail to capture the nonideality of the imaging process and the full statistical structure of microscopy images. Hence, *data-driven* approaches can be more broadly applicable and better suited than analytical ones to solve image reconstruction problems. Ideally, we want generic reconstruction algorithms that can be trained from exemplary data obtained either experimentally or from simulation. This is how Deep Learning comes into play. Deep neural networks are capable of learning end-to-end image transformations from data without the need for explicit modelling (see **Box 3**).

## Fluorescence microscopy image reconstruction.

Image reconstruction in fluorescence microscopy encompasses a wide range of topics ranging from image denoising, deconvolution, registration, stitching, fusion, super-resolution, as well as challenges brought by more exotic acquisition modalities such as tomographic, light-field, lens-less, and deep-tissue imaging. Here, we review existing applications of Deep Learning in fluorescence microscopy and suggest promising new ones.

**Image Denoising.** Image noise in fluorescence imaging is typically caused by the combined effect of Poisson statistics of photon-counting and imperfect electronics of the sensor. Early on, denoising algorithms were recognised as an opportunity to speed-up acquisition by using shorter exposures and reduce photo-damage by decreasing illumination intensity<sup>31</sup>. Algorithms were specially developed for fluorescence images that leveraged compressed sensing<sup>32</sup>, image self-similarity<sup>31,33</sup> and thresholding after wavelet transform<sup>34</sup>. Today, state-of-the-art

\* Correspondence: loic.royer@czbiohub.org



**BOX 1** The Deep Learning revolution.

In 2012 Krizhevsky, Sutskever and Hinton published their now famous AlexNet paper<sup>17</sup>. In that work, they showed that a deep convolutional neural network (CNN) could win the *ImageNet Large-Scale Visual Recognition Challenge* (ILSVRC) by an unprecedented margin. AlexNet was trained on 1.2 million high-resolution training examples to classify images into a thousand object categories. The double digit decrease in error rate achieved by Hinton's team astounded the computer vision community. One year later, most competing teams used CNNs for what was considered

at the time the most challenging machine learning competition. Since 2012, deep CNNs have continued to break records going from an error rate of 15.4%<sup>17</sup> in 2012 to just 3.08%<sup>18</sup> in 2018 – surpassing human accuracy. Nowadays, deep CNNs are the state-of-the-art for major machine learning challenges on object<sup>18</sup>, and speech recognition<sup>19</sup>. Recent striking Artificial Intelligence successes critically depend on deep CNNs, e.g. DeepMind's AlphaGo victory against the second best human GO player<sup>20</sup>. This has contributed to the enthusiasm, but also hype, surrounding Deep Learning.

classical denoising algorithms in computer vision research use sparse coding techniques<sup>35</sup>, low-rank decomposition<sup>36</sup>, or image self-similarity<sup>37,38</sup>. The lesson learned over time is that to effectively denoise an image, it is better to look beyond small patches of pixels and gather information across multiple image regions and ideally *learn* from a large sets of images. In 2009 Jain et al.<sup>39</sup> trained a 5-layer CNN to denoise natural images achieving a lower error rate and two orders of magnitude faster processing than previous approaches. Since then, numerous deep CNN architectures have been proposed for denoising natural images<sup>40</sup>. Recently, Weigert et al.<sup>8</sup> trained networks with low-SNR (input) and high-SNR (target) image pairs to demonstrate effective denoising of fluorescence microscopy images acquired under low light or short camera exposure conditions (see Fig. 1a). They found that their method maintains image quality while decreasing light exposure 60-fold and outperforms classical denoising such as NLM<sup>37</sup> or BM3D<sup>38</sup>. Such Deep-Learning-based restoration can be useful when specimens are light-sensitive, fluorophore bleaching is of concern, or to study fast dynamics in live samples. A major obstacle to applying Deep Learning for image denoising is obtaining matched pairs of high-noise and low-noise training images. This requirement can be relaxed as shown by recent developments in unsupervised and self-supervised image restoration such as *noise2noise*<sup>41,42</sup>, *noise2self*<sup>43</sup>, *noise2void*<sup>44,45</sup>, *deep image prior*<sup>46</sup>, or novel unpaired image transformation schemes such as *cycleGAN*<sup>47</sup> (see Fig. 2h). We expect that these advances will widen applicability and facilitate adoption of CNN-based denoisers in fluorescence microscopy.

**Spatial Deconvolution.** The resolution of images obtained by fluorescence microscopy is fundamentally limited by the optical parameters of the microscope. Image deconvolution can improve image quality by enhancing high-frequency details. In the case of 3D fluorescence microscopy this is particularly useful for reducing axial blur and out-of-focus light<sup>48,49</sup>. Classic Bayesian approaches such as Richardson-Lucy deconvolution<sup>21</sup> are still used extensively and have been successfully extended to multi-view light-sheet microscopy<sup>50</sup>. A pervasive problem in microscopy is poor axial sampling and resolution that leads to 3D image anisotropy – i.e. lower resolution along the ob-

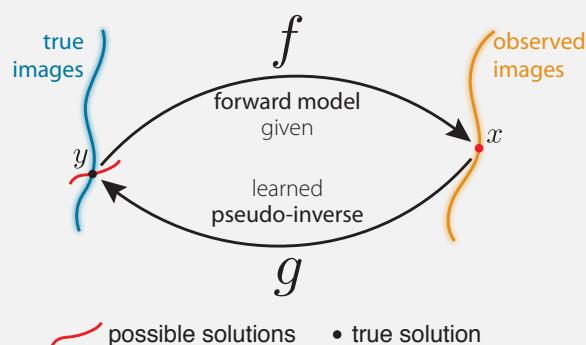
jective axis (z) than in the transverse (x-y) planes. To address this, Weigert et. al<sup>8,9</sup> recently demonstrated a self-supervised training strategy that can deconvolve axial (xz) images and restore isotropic resolution, significantly outperforming the popular Richardson Lucy algorithm. Challenges remain: when imaging large samples, the point-spread-function (PSF) varies in space and across views and even possibly over time. Shajkofci et al. recently showed how a CNN can be used to estimate optical aberrations and improve image quality by 1dB compared to other deconvolution algorithms<sup>10</sup>. It would be interesting to investigate whether the same approach can be applied to volumetric imaging. Another interesting opportunity would be to address challenges that are opto-mechanical in nature: in light-sheet microscopy there is a limited choice of illumination and detection objectives that can be placed orthogonal to each other. To solve this problem one could use lower numerical aperture (NA) objectives that mechanically fit, and apply Deep Learning to computationally increase the detection NA (see Rivenson et al.<sup>51</sup>).

**Super-resolution.** PALM<sup>52</sup> and STORM<sup>53</sup> achieve super-resolution by localising sparse fluorophore emissions in thousands of images. These images are an intermediate product that must be processed to reconstruct a super-resolved image of fluorophore locations<sup>54</sup>. In a technique called Deep-STORM<sup>13</sup>, Nehme et al. have demonstrated the use of an encoder-decoder network (see **Box 4**) to localise emitters (see Fig. 1b). Interestingly, the authors used simulated pairs of diffraction-limited and super-resolved images to train their network and obtain high emitter detection efficiency of 88% on simulated images, and as high as 96% when trained with actual experimental data. Instead of predicting high-resolution images, Boyd et al.<sup>55</sup> demonstrated direct prediction of the coordinates of emitters. In another application of Deep Learning to super-resolution microscopy, Ouyang et al.<sup>12</sup> trained a U-Net (see **Box 4**) fed with sparsely sampled PALM images to predict their densely sampled counterparts (see Fig. 1c). In contrast with the mean-squared-error (MSE) training loss of Deep-STORM, their ANNA-PALM method uses an ingenious combination of  $L_1$ , structural-similarity (SSIM), and conditional adversarial losses<sup>29,30</sup> to obtain good quality image reconstruction. Us-

## BOX 2 Image reconstruction, inverse problems, and Deep Learning.

In most imaging modalities we have a well founded understanding of the physics behind the image formation process. We can write formulas that describe, and code that simulates, how light propagates through optical elements to eventually form an image on a detector – this is called the ‘forward model’  $f$ . This model is a transformation from the desired *ideal* image  $y$  to *observed* images  $x$ . Observed images  $x$  are incomplete, degraded, or convoluted compared to  $y$ . In the opposite direction, reconstructing the true image  $y$  from the observed images  $x$  is a challenging *inverse problem*. Finding the inverse of  $x$  necessitates a pseudo-inverse function  $g$  that is classically implemented as an iterative algorithm. One obstacle to inversion is that forward processes are often stochastic. Another fundamental difficulty is that the forward model maps many images to the same observation, and so there is no well-defined inverse but instead multiple confounding solutions (red line). However, in most cases, we have *a priori* knowledge about validity of solutions, which makes it possible to select one inverse. Until recently, this *prior* information about the solution had to be mathematically formulated by experts and engineered into algorithms. These handcrafted priors typically prescribe, for example, that correct solutions should be smooth (Tikhonov’s and total variation regularisations) or admit a concise representation in some basis (wavelet

basis, dictionary) <sup>22</sup>. However, the statistical structure of real microscopy images is far more complex – limiting the quality of images reconstructed with such priors. Hence the advent of *data-driven* inversion schemes that learn a pseudo-inverse  $f$  from large datasets of  $(x, y)$  pairs. These pairs can be obtained by simulation – i.e. by computing  $f$ , or can be obtained empirically by means of clever experimental strategies that produce both  $x$  and  $y$ . The learning procedure not only learns the function  $g$  but also, implicitly, the prior knowledge about ‘good’ solutions. Deep Learning is one promising data-driven approach for solving inverse problems, and by extension image reconstruction tasks<sup>23</sup>.



ing a similar network architecture and training loss, Wang et al.<sup>11</sup> have recently demonstrated that the resolution of diffraction limited confocal images can be enhanced (see Fig. 1d). This work shows that use of structured losses – whether explicit (SSIM) or adversarially learned (cGAN, see Box 4) – result in output images that are sharper and of better perceptual quality<sup>30</sup>. In order to perform super-resolved fluorescent imaging on live cells, the sparsity requirement must be abandoned which has led to the development of algorithms capable of handling higher emission densities – such as SRRF<sup>56</sup>. We expect that the use of CNNs in this challenging regime could have a significant impact in further improving image quality and temporal resolution.

**Structured illumination.** Another super-resolution approach is structured illumination microscopy (SIM). It can surpass the diffraction limit 2-fold by computational synthesis of images acquired by shifted illumination patterns<sup>57,58</sup>. However, current algorithms require precise knowledge of the effective pattern and are thus sensitive to distortions or attenuation caused by the sample. While classical algorithms exist that can handle unknown<sup>59</sup>, or distorted<sup>60</sup> illumination patterns; it would be interesting to compare these with a CNN-based SIM reconstruction algorithm (see Fig. 2c) possibly leveraging ideas similar to magnetic resonance imaging reconstruction with AUTOMAP<sup>7</sup>. The training data could be generated by simulations that would

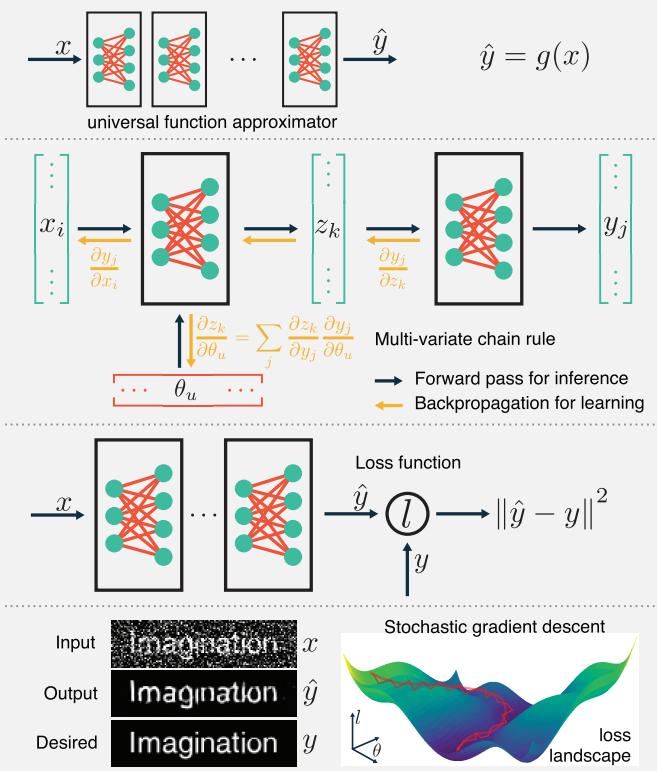
take into account all relevant degradations and distortions. We believe that such an algorithm could help obtain better SIM reconstructions under challenging signal-to-noise conditions.

**Spectral Deconvolution.** The trend in fluorescence imaging is for faster acquisition with more colours. However, fluorescent proteins and dyes used for labelling have broad overlapping emission spectra – which leads to channel mixing that complicates interpretation. One solution that sacrifices speed is to acquire the channels sequentially. Another solution is to acquire the channels simultaneously – but that requires application of spectral unmixing algorithms for channel separation. To address this problem, algorithms such as linear unmixing<sup>61</sup> or phasor approaches<sup>62</sup> have been devised to *deconvolve* these spectra into separate channels corresponding to distinct labels. In some sense, unmixing multi-color acquisitions is a deconvolution problem in the spectral dimension that is difficult in low signal-to-noise conditions. Could CNNs be applied to deconvolve more wavelengths in poor signal-to-noise conditions? One straight-forward approach to obtain training data would be to acquire each channel sequentially and then simultaneously (see Fig. 2b), or use physics-based models to generate synthetic training data.

### BOX 3 Deep Learning as Differential Learning

In Deep Learning non-linear parameterised processing modules are combined to progressively transform an input  $x$  into the desired output  $y$ . The adjective *deep* refers to the large number of stacked modules required for this challenging task. The learning process begins with a random choice for the parameters  $\theta$  of the modules, which are iteratively updated to improve the accuracy of the reconstruction. This is possible because each module is *differentiable*, meaning that we know exactly how changes in parameter values cause changes in output values. More formally, we know the partial derivatives  $\frac{\partial \hat{y}}{\partial \theta}$  of the output  $\hat{y}$  with respect to any parameter  $\theta$ . Artificial neural networks are the most popular implementation of such modules. Networks with a sufficient number of parameters, and at least three layers, can theoretically approximate any function<sup>24</sup>.

The backpropagation algorithm<sup>25</sup> uses the chain-rule to efficiently compute all partial derivatives – or gradients – with just one forward pass through the network followed by a backward pass. The discrepancy between the desired output  $y$  and the actual output  $\hat{y}$  is quantified using a loss function  $l(\hat{y}, y)$ . The parameters  $\theta$  are updated to minimise this loss using *stochastic gradient descent*. The speed provided by computing the updates on small batches of data, in parallel, on GPUs, allows one to fit networks with millions of parameters on datasets with millions of observations.



**Correcting uneven illumination and detection.** Unfortunately, in most microscopes, illumination and detection efficiency are not uniform over the entire field of view. In wide-field or spinning disk confocal microscopes<sup>63</sup>, uneven detection efficiency often requires non-uniformity correction. In the case of mesoscale light-sheet microscopy, the problem is far more complex because illumination shadows are caused by sample-induced occlusion<sup>64</sup>. We believe that CNNs are well positioned to address this problem. Promising CNN-based techniques exist in computer vision for context-based *inpainting* of missing image regions<sup>65,66</sup>. Fortunately, in lightsheet microscopy information is not completely absent because the occlusion is often times only partial. Moreover, in the case of multi-view imaging, training data is readily available in the form of images acquired from different detection arms or with light-sheets of distinct propagation angles.

**Scattered light haze reduction.** Scattered light is a significant cause of image quality deterioration in live fluorescence imaging. For example, in lightsheet microscopy, background induced by scattered light is often suppressed to facilitate subsequent image processing steps such as image stitching and fusion<sup>67</sup>. Yet, many challenges remain because scattering is depth-sensitive and sample dependent. Recently, powerful hybrid approaches have been developed that combine CNNs and optical modelling for dehazing natural scenes<sup>68</sup>. Could such

ideas be applied to microscopy to further improve the quality of images? In multi-view light-sheet microscopy an obvious way to obtain training data is to collect low/high-haze image pairs from overlapping regions obtained from distinct views (see Fig. 2a). In some cases, matching image pairs cannot be obtained, such as for example when imaging deep within a sample using a spinning-disk confocal microscope. Instead, recent advances in unpaired image-to-image translation<sup>47</sup> could be used to learn to *dehaze* images using images at the surface and deeper within the sample (see Fig. 2h).

**Registration, stitching and fusion.** Because of the limited field-of-view of objectives and large size of some specimens, it is often necessary to break-down an acquisition into 2D or 3D tiles<sup>50</sup>. Reconstructing a complete image requires stitching the tiles together after registration to a common reference frame. State-of-the-art algorithms<sup>50,69</sup> rely on iterative schemes that leverage direct image cross-correlation or matching fiducial markers. One obvious way to apply Deep Learning to image registration is to teach CNNs to find image correspondences<sup>70</sup>, and thus have a Deep Learning component be part of a larger algorithm. Yet, one key promise is that a single network can be trained to solve a complex multi-step task in an *end-to-end* fashion. Recently, Rohe et al.<sup>71</sup> showed that pairs of large 3D medical images can be non-rigidly registered in less than 30ms using a fully convolutional CNN that directly outputs a defor-

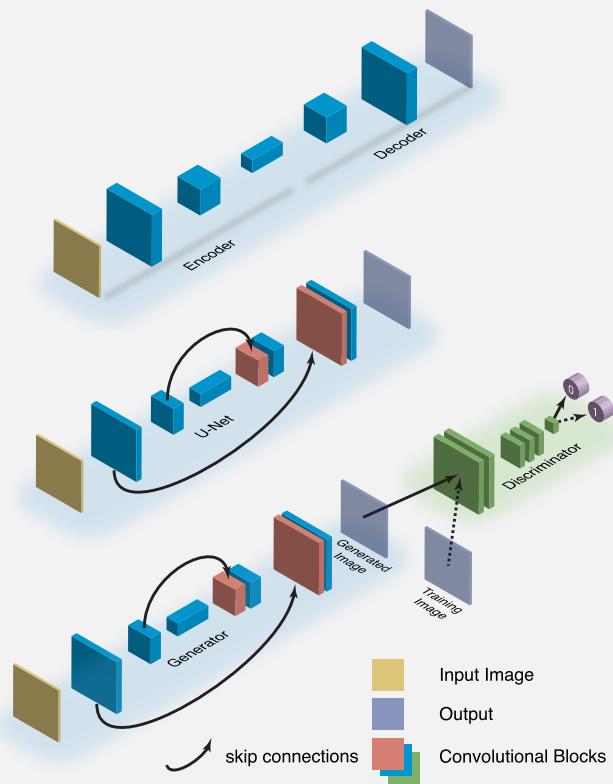
#### BOX 4 Architectures for image reconstruction

Image restoration or reconstruction require neural networks that map one or several input images to an output image. Input and output images may be 2D or 3D and may even have different dimensions. The encoder-decoder convolutional network architecture<sup>26</sup> is one of the simplest networks capable of image-to-image translation. First, the input image is successively and repeatedly downsampled and convolved – in addition to other normalisation and non-linear operations typical for CNNs. At every step, the image dimensions are reduced while the number of image *channels* – an additional dimension – is increased. This creates a *representation bottleneck* that is then reversed with successive and repeated upsampling and convolution steps. The bottleneck forces the network to *encode* the image into a small set of abstracted variables (often called *latent variables*) along the channel dimension which are then *decoded* by the second half of the network. This ensures that only the most relevant features from the input images are retained and those that are deemed unessential to the reconstruction are discarded.

The U-Net<sup>27</sup> is a variation of the encoder-decoder network. It adds *shortcut* or *skip* connections between the encoder and decoder branches. These additional connections are beneficial for image restoration tasks because they send fine image details directly to the decoder pass – skipping the bottleneck. For image restoration tasks where the identity function is a good first guess, additive skip connections can be used to learn a residual mapping. Overall, skip and residual connections have been shown to help train deeper networks by preventing vanishing gradients<sup>28</sup>.

In the generative adversarial network (GAN)<sup>29</sup> framework, two networks, a generator and a discriminator, are trained together. The generator learns to generate images that mimic examples from the training set while the discriminator learns to classify images as real (from the training set; label 1) or counterfeit (from the gener-

ator; label 0). In practise, this behaviour is achieved by training the generator to minimise the loss function and training the discriminator to maximise it. At the end of training, the discriminator network is dropped and only the generator is used for image reconstruction. A recent improvement on GANs are conditional GANs (cGAN) which also provide the input image to the discriminator<sup>30</sup>. cGANs can be understood as an implicitly learned loss function defined by the action of the discriminator.



mation vector field. Their method is more accurate and faster than traditional approaches, does not require iterations of an optimiser, but does require a large amount of training data. Similarly, Nguyen et al.<sup>72</sup> demonstrated an *unsupervised* approach for non-rigid transformation estimation that is faster and more accurate than classical correlation-based or fiducial-based approaches. We expect that these ideas will eventually percolate into the field of fluorescence imaging. An upcoming and exciting challenge is imaging of live moving animals. For example, *C. elegans* larvae twist incessantly before hatching making the analysis of time-lapse data challenging. This requires specialised stabilisation algorithms capable of unfolding, straightening and registration<sup>73</sup>. Can a CNN learn to stabilise 3D time-lapses images of highly dynamic samples against a fixed template shape (see Fig. 2d)?

**Light-field, lens-less imaging, and tomography.** Recently, exotic forms of fluorescence imaging have emerged that pose novel computational challenges. For example, light-field imaging can multiplex an entire volume into a single 2D image acquisition thus enabling rapid functional imaging at single neuron resolution<sup>74</sup>. Recent work by Fei et al.<sup>75</sup> demonstrated some initial results on CNN-based light-field 3D image reconstruction of live *C. elegans* worms. Similarly, lens-less imaging schemes such as *DiffuserCam* achieve one-shot 3D imaging by replacing the objective lens with a diffusing element placed in front of the sensor<sup>76</sup>. Tomography is another indirect approach to obtain 3D images from 2D acquisitions that is particularly relevant for mesoscopic fluorescence imaging<sup>77</sup>. In both cases the final 3D image is computationally reconstructed from a single 2D image. We believe that both lens-less imaging and tomography will benefit from Deep-Learning-based algorithms –

just as it has been demonstrated for light-field imaging. Indeed, recent results give strong evidence that CNN-based approaches are superior to classical handcrafted techniques for reconstructing 3D images from 2D acquisitions<sup>7</sup>.

**Temporal consistency.** Applications of Deep Learning to microscopy typically fail to capture temporal patterns: predictions are made independently, frame-by-frame. As a consequence, successive reconstructed images are not necessarily *temporally consistent* and can exhibit artefacts that are clearly noticeable when examined across time. Explicitly modelling the temporal dimension and training with time-resolved data would lead to more accurate reconstructions. A straightforward approach for augmenting any CNN with temporal consistency is simply to treat time as an additional dimension (see Fig. 2f). However, this can become intractable for large networks handling long-term correlations. Instead, a better alternative is to combine CNNs with state-of-the-art recurrent neural networks which have been specifically designed for sequence prediction. One obvious candidate is the *convolutional Long-Short Term Memory* (convLSTM) architecture which combine CNNs for space and LSTMs for time and have been successfully applied to weather forecasting<sup>78</sup>.

**Beyond fluorescence.** In some cases, fluorescence measurements are a proxy for other quantities. For example, measuring flow by fluorescence is an important tool in the emerging field of cell mechanobiology<sup>79</sup>. Unfortunately, current hand-crafted *optical flow* algorithms typically fail to produce smooth flow fields in challenging imaging conditions. CNN-based optical flow estimation algorithms recently developed for self-driving cars<sup>80</sup> could be a source of inspiration for designing more robust algorithms (see Fig. 2e). Another interesting challenge is registering electron microscopy (EM) and fluorescence microscopy images. This task often requires placement of fiducial markers visible in both modalities<sup>81</sup>. Recent work on virtual fluorescence labelling<sup>14–16</sup> suggests that it could be possible to automatically translate images across microscopy modalities. Unsupervised approaches for learning image correspondence such as cycleGAN<sup>47</sup> could be a powerful approach in correlative electron-fluorescence imaging. A CNN could be trained to translate between electron and fluorescence microscopy images from unregistered pairs. Fiducial markers could then be localised between EM and fluorescence images for image registration.

**Deep Learning for adaptive and smart imaging.** State-of-the-art microscopes – in particular light-sheet microscopes – are becoming adaptive machines capable of closed-loop control and real-time image analysis<sup>82,83</sup>. Yet, current adaptive microscopes rely on time-consuming and complex iterative schemes<sup>82</sup> or direct measurements<sup>83</sup>. What if one could design a system that quickly and correctly guesses the best imaging parameters from few measurements? This is an inverse problem, because it may be easy to deduce the consequence of a particular choice of parameters, but difficult to determine which choice

will lead to the desired image. There exist already some limited examples that show, for example, how to train a CNN to focus light deep within turbid media<sup>84</sup>. Yet, the true challenge going forward will be to optimise multi-step imaging strategies: the microscope will observe, make *moves* that change imaging parameters, and thus play a *game* with the ultimate goal of gaining the most information about the sample (see Fig. 2i). Deep reinforcement learning – another key concept behind DeepMind’s AlphaGo success<sup>20</sup> (see **Box 1**) – could be used to automatically learn such strategies. Ultimately, we expect that most of the human decision making required for optimising imaging will be done automatically – something that will also necessarily require extensive end-to-end robotic automation so that the ‘game’ can be restarted at will.

### Limits, challenges and pitfalls

Despite its successes, Deep Learning also comes with risks, has limits, and brings new challenges.

**The hallucination problem.** The most serious issue when applying Deep Learning for discovery is that of hallucination. When looking at random patterns such as clouds, our brains can perceive shapes of objects and animals. Similarly, Deep Learning systems can hallucinate details and make mistakes when provided with inadequate training data (see **Box 5**). Some network architectures – such as GANs – are particularly susceptible to hallucination because their explicit goal is to fool a discriminator by forging persuasive details. Conditional adversarial networks (cGANs, see **Box 4**) as used by Isola et al.<sup>30</sup> still produce plausible image details that are not necessarily supported by the input images, which is a potential concern for using GANs or cGANs for discovery. A possible mitigation strategy is to add additional consistency losses to prevent hallucinations as done in ANNA-PALM<sup>12</sup>. In general, mistakes made by neural networks can be highly plausible and more subtle than those of classical algorithms.

**The over-fitting problem.** Another failure mode of neural networks is over-learning or over-fitting: the neural network memorises exactly the training data and loses the ability to generalise to unseen data. This can be caused by insufficient training data, or by a poor choice of network parameters. Architectural features such as drop-out regularisation can help, but the best way to avoid over-fitting is to train with data that exceeds the memorisation capacity of the networks. Yet, in practice, even if networks do not overfit *per-se*, they do remain fragile to subtle variations between training and application data<sup>85</sup>. Moreover, imperceptible image variations can be engineered to confuse neural networks<sup>86</sup>. Overall, current Deep Learning approaches cannot generalise well to data obtained from a different microscope or under different conditions<sup>8</sup>.

**Limits imposed by quality and availability of training data.** The success of Deep Learning depends critically on availability of training data. If the amount of training data is not suffi-

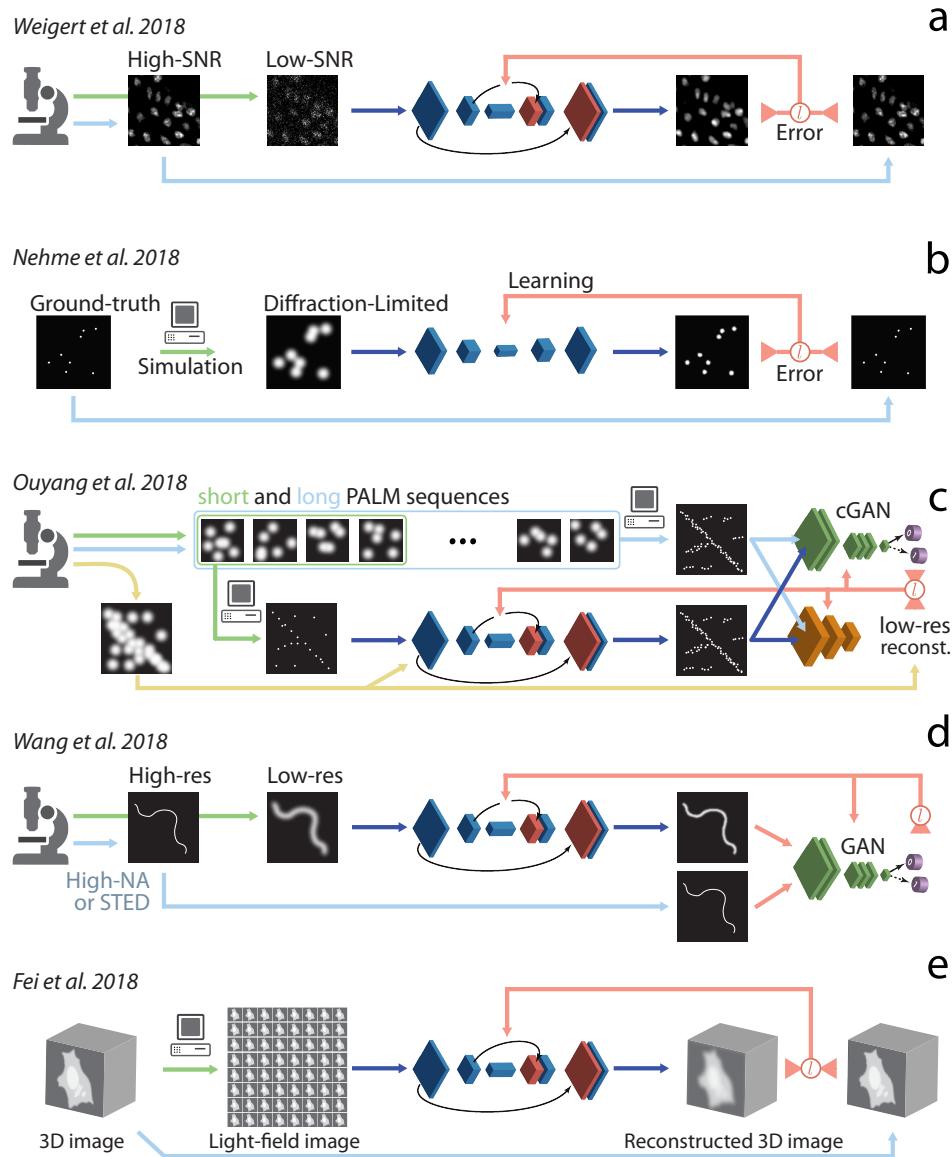


Figure 1: Review of current applications of Deep Learning to fluorescence microscopy. **(a)** Content-aware image restoration (CARE) for denoising, super-resolution and axial deconvolution<sup>8,9</sup>. **(b)** DeepStorm for STORM super-resolution<sup>13</sup>. **(c)** AN-NAPALM for accelerated PALM super-resolution<sup>12</sup>. **(d)** Computational fluorescence super-resolution by learning to translate from low to high numerical aperture (NA) images<sup>11</sup>. **(e)** Fluorescence light-field reconstruction for high-speed high-resolution volumetric imaging<sup>75</sup>.

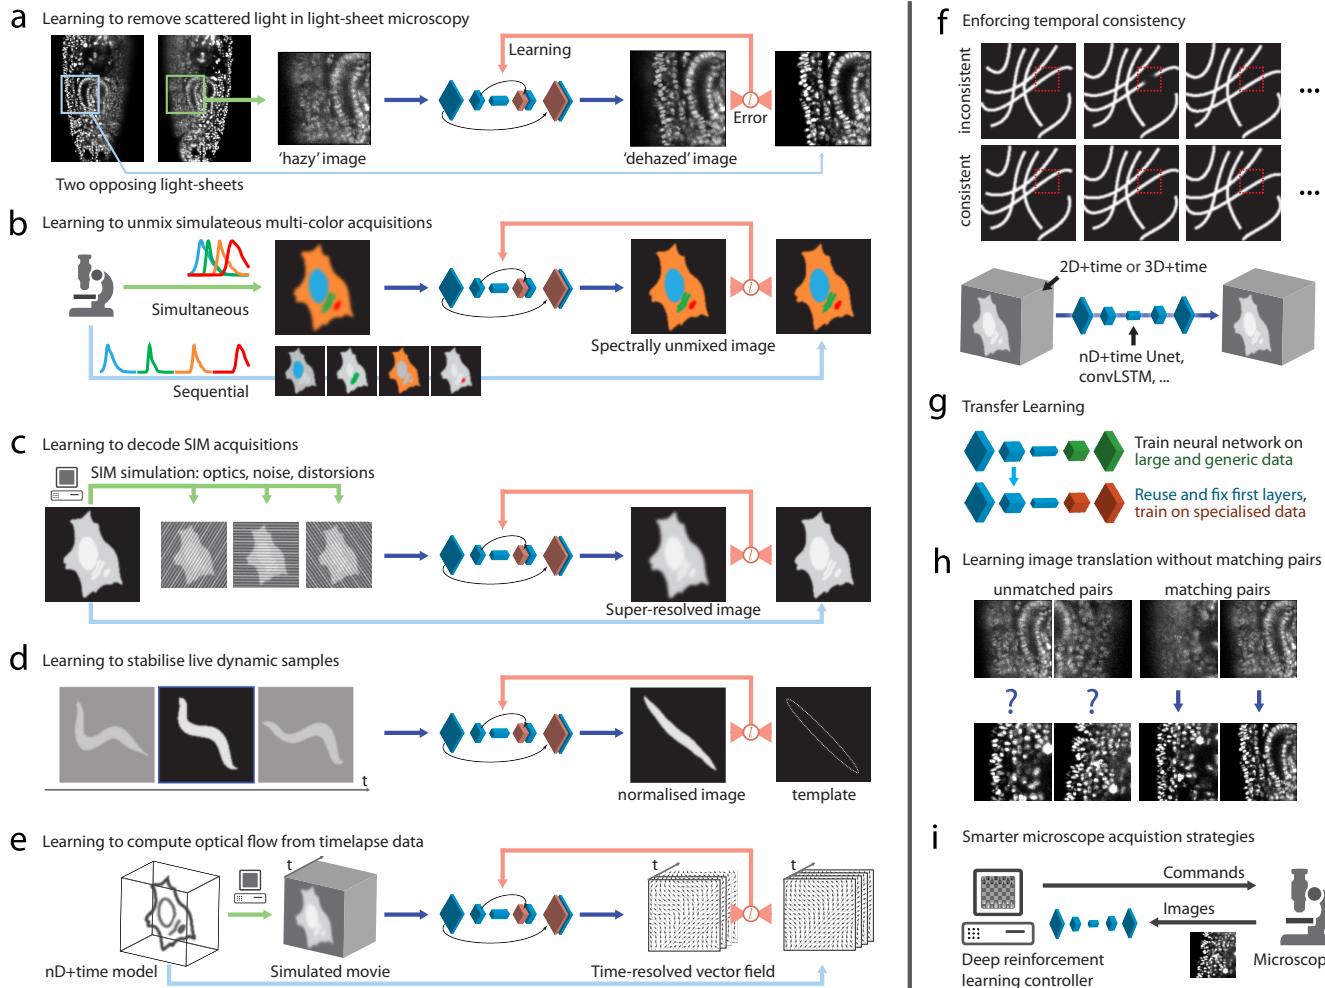


Figure 2: Potential applications of Deep Learning in fluorescence microscopy and key concepts. **(a)** Learning to reduce scattered light 'haze' in light-sheet microscopy. **(b)** Learning spectral unmixing of simultaneous multi-color acquisitions. **(c)** Learning to reconstruct super-resolved images from structured illumination acquisitions. **(d)** Learning to straighten live dynamic samples with weak supervision against a template shape. **(e)** Fluorescence optical flow with Deep Learning. Forward model simulations can produce time-lapse data for given vector fields, and a network can be trained to compute the inverse: vector fields from fluorescence spatio-temporal fluctuations. **(f)** Leveraging the time dimension by enforcing temporal consistency in predictions to further improve image quality – applicable to any image reconstruction task. **(g)** Transfer learning and other approaches to facilitate training and reducing the need of specialised curated training datasets will be key for the success of Deep Learning in microscopy. **(h)** Unsupervised learning is another research direction that will reduce the need for curation. For example, unpaired image translation (CycleGANs<sup>47</sup>) are a promising approach. **(i)** Deep Reinforcement Learning could drive smart microscopes able to conduct complex imaging experiments with the goal of acquiring the best possible image.

## BOX 5 Deep Learning based discovery, and its limits.

Does data-driven image reconstruction preclude discovery? Can novel structures and unknown patterns be discovered even if absent from the training data? To shed some light on this subtle but important question we will use an analogy: let us train a neural network (U-Net) to restore a highly degraded image of an ancient English word: 'Witenagemot'. First, we use a training dataset consisting of images of all 3 letter sub-strings occurring in the most common English words. Importantly, we exclude 'Witenagemot' and its variants. As shown, the word is successfully recovered. We have just *discovered* a word that is not present in the training data. Knowing which alphabet to expect does not preclude the discovery of new words, sentences or ideas. What if we exclude the letters 'a' and 'e' from the training data? In that case, the network fails to recognise these letters, and discovery of the word is hampered. Now, if one uses the wrong alphabet for training, i.e. 6000 Chinese characters, discovery is now impossible. The mistakes made are interesting in their own right: the network tries to reconcile the inadequate prior with

the information present in the degraded image – leading to creative hallucinations. Finally, we can rescue the restoration by using instead a 50%-50% mix of Chinese and Latin characters. However, we also see that some characters like the 'm' are decorated with artefacts of Chinese provenance.



cient, poor performance will ensue (see **Box 5**). Yet, a common misconception is that Deep Learning *always* requires a *very large* number of training examples. More accurately, the needed amount really depends on the problem difficulty. Far less training data is required in fluorescence microscopy than is needed to distinguish between, for example, 800 different kinds of birds in natural images (ImageNet challenge<sup>17</sup>). Indeed, in some rare cases, the training and inference can be done on the very same image (see Isonet<sup>9</sup>). In fact, as shown in **Box 5**, the quality of the data and its suitability for the problem is perhaps more important. In any case, fluorescence microscopy needs creative experimental and computational strategies to obtain more and better training data.

**Challenges in obtaining training data.** A classic approach to increase the number of training examples is *data augmentation* which creates variants of existing images through image rotation, scaling, lighting, etc. In cases where the physics of the image degradation process is well understood, forward-model simulations can also be used to generate training data<sup>8,13</sup>. Interestingly, neural networks could also be used to improve the quality of these simulations. For example, Osokin et al.<sup>87</sup> trained a GAN model to generate realistic images of fluorescently tagged yeast cells – which could in turn be used for training restoration or reconstruction algorithms. Another popular approach is *transfer learning* that has been shown to accelerate convergence and improve generalisation by pre-training networks on distinct and large datasets<sup>88</sup> (see Fig.2g). It relies on the observation that the first layers of a neural network function as detectors of universally applicable textures and patterns<sup>89</sup>. As an example, Estava et al.<sup>90</sup> recently demonstrated how Google's Inception network trained on ImageNet

images can be re-trained to classify dermatological images for melanoma classification. Instead of using ImageNet, could we use existing large collections of fluorescence microscopy images? Prime candidates are the cellular and organelle fluorescence images from the *Human Protein Atlas*<sup>91</sup> or the *Broad Bioimage Benchmark Collection*<sup>92</sup>. Sullivan et al.<sup>93</sup> combined crowd-sourced annotations with machine learning on an image-classification task – could some of these ideas be applied to fluorescence image reconstruction? Or, at least, could such repositories be used for transfer learning?

**Interpreting and trusting the black box.** Another crucial challenge is that deep neural networks are black-box models whose operation is opaque and difficult to interpret. For Deep Learning to become a trusted component of the microscopy-based discovery pipeline, we need conceptual frameworks, and interactive graphical tools, to help explain how a particular result is attained. Fortunately, progress has been made within the computer vision community in this direction with techniques developed to: (a) better understand which parts of the input image are essential to predicting the output<sup>94</sup>, (b) analyse the function of intermediate layers<sup>89</sup>, (c) tease apart contributions from different network layers with *ablation studies*<sup>89</sup>, (d) build hierarchical explanatory graphs across layers<sup>95</sup>, and (e) design network architectures that are interpretable by design<sup>96</sup>. These ideas need to be adapted to fluorescence imaging and tools should be built to facilitate the interpretation of results. One possibility is to devise tools that would provide explanations for given predictions. Take for example, low-contrast images of fluorescently labelled cell membranes that would be restored using a CNN. The user could select pixels along a predicted membrane and ask for an explanation – for example in the form

of a covariance map – that could be presented interactively to show which pixels and patterns in the input image justify the reconstruction and why. Yet, it is not enough to explain how a result is attained, it is also necessary to quantify the confidence in these results. Recently, several methods for computing pixel-wise confidence measures and confidence intervals have been proposed that leverage ensemble of predictors and adversarial training<sup>8,97</sup>. Another promising approach is to utilise internal consistency checks – for example by comparing the input of the network to the output transformed by the forward model as done during inference in ANNA-PALM<sup>12</sup>.

**Is discovery possible?** An oft-asked question is: “Is it possible to reconstruct structures or patterns that are not present in the training data?”. To illustrate this problem, we use the restoration of degraded text as an analogy for image reconstruction in microscopy – see **Box 5**. What are the lessons to be learned? First, data-driven discovery is possible, but it depends critically on the quality and compatibility of the training data. There are also serious hallucination risks when the training data contains less or more information than strictly required. The same is true for microscopy: the recognition of image features does not preclude further analysis of their spatial organisation nor the study of their temporal dynamics. For example, being able to super-resolve microtubules with fewer raw image acquisitions makes it possible to make observations at higher temporal resolution, and consequently opens the door to new discoveries<sup>12,13</sup>. And yet, mistakes can be made if the training data is inadequate. In light of these results, it is our opinion that Deep Learning results cannot be blindly trusted – something certainly true of any singular piece of evidence. While confidence measures are needed<sup>8</sup>, discoveries must be confirmed from multiple lines of evidence gathered from distinct experiments.

**Reproducibility.** Reproducibility is the bedrock of modern science and it is only natural to expect independent validation of published Deep Learning results. However, a lack of experimental consistency and reporting guidelines sometimes prevents Deep Learning researchers from replicating each other’s results<sup>98</sup>. Moreover, the statistical significance of measured performance metrics is rarely reported in the published literature because of high cost (in terms of time and resources) associated with training<sup>99</sup>. To ameliorate this, researchers should adopt a set of guidelines when publishing their results. Such best practices could include, for example: (i) making the source code and trained models freely available, (ii) allowing access to the dataset that was used to train the models, (iii) full disclosure of all hyper-parameters used during training, (iv) reporting the statistical significance of observed results, and (v) reporting failure cases in addition to exemplary ones. Such approaches will increase confidence in the results and facilitate widespread adoption of Deep Learning techniques in fluorescence microscopy.

### Dissemination.

In order to accelerate adoption of Deep Learning in microscopy, novel software frameworks tailored for biologists are needed for using, adapting, training, validating and interpreting deep neural networks<sup>8</sup>. In particular, tools with better ergonomics, user-friendly graphical interfaces, and smart parameter-free or auto-tuning algorithms would help lower the technical bar to adoption. Another potential obstacle is the cost of hardware for training. While a custom-built desktop computer outfitted with GPUs is relatively affordable, assembling and maintaining such a machine requires significant technical skills. An alternative is pay-as-you-go cloud-based Deep Learning platforms which are best reserved for short-term high-intensity workloads.

### Conclusion

Deep Learning holds many promises for fluorescence microscopy. Some will stand the test of time, others will fall short. Some applications will be redefined by it, others will continue to rely on simpler methods. In either case, *learned algorithms* will be key for upcoming computational advances in microscopy. In future, we expect that Deep Learning models trained end-to-end will blur the frontier between image reconstruction and image analysis. At present, input and outputs are images, but in future we foresee models capable of transforming images directly into knowledge.

### Bibliography

1. Lichtman, J. W. & Conchello, J.-A. Fluorescence microscopy. *Nature methods* **2**, 910–919 (2005).
2. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436 (2015).
3. Lam, C. & Kipping, D. A machine learns to predict the stability of circumbinary planets. *Monthly Notices of the Royal Astronomical Society* **476**, 5692–5697 (2018).
4. Webb, S. Deep learning for biology. *Nature* **554**, 555–557 (2018).
5. Radovic, A. *et al.* Machine learning at the energy and intensity frontiers of particle physics. *Nature* **560**, 41–48 (2018).
6. Xu, Y. *et al.* Deep learning of feature representation with multiple instance learning for medical image analysis (2014).
7. Zhu, B., Liu, J. Z., Cauley, S. F., Rosen, B. R. & Rosen, M. S. Image reconstruction by domain-transform manifold learning. *Nature* **555**, 487–492 (2018).
8. Weigert, M. *et al.* Content-Aware Image Restoration: Pushing the Limits of Fluorescence Microscopy (2017).
9. Weigert, M., Royer, L., Jug, F. & Myers, G. Isotropic Reconstruction of 3d Fluorescence Microscopy Images Using Convolutional Neural Networks. In *Medical Image*

*Computing and Computer-Assisted Intervention – MIC-CAI 2017*, 126–134 (Springer International Publishing, 2017).

10. Shajkofci, A. & Liebling, M. Semi-blind spatially-variant deconvolution in optical microscopy with local point spread function estimation by use of convolutional neural networks. *arXiv preprint arXiv:1803.07452* (2018).
11. Wang, H. *et al.* Deep learning achieves super-resolution in fluorescence microscopy. *bioRxiv* 309641 (2018).
12. Ouyang, W., Aristov, A., Lelek, M., Hao, X. & Zimmer, C. Deep learning massively accelerates super-resolution localization microscopy. *Nat. Biotechnol.* **36**, 460–468 (2018).
13. Nehme, E., Weiss, L. E., Michaeli, T. & Shechtman, Y. Deep-STORM: super-resolution single-molecule microscopy by deep learning. *Optica, OPTICA* **5**, 458–464 (2018).
14. Christiansen, E. M. *et al.* In Silico Labeling: Predicting Fluorescent Labels in Unlabeled Images. *Cell* **173**, 792–803.e19 (2018).
15. Ounkomol, C., Seshamani, S., Maleckar, M. M., Collman, F. & Johnson, G. R. Label-free prediction of three-dimensional fluorescence images from transmitted-light microscopy. *Nature methods* **15**, 917–920 (2018).
16. Rivenson, Y. *et al.* Deep learning-based virtual histology staining using auto-fluorescence of label-free tissue <http://arxiv.org/abs/1803.11293v1>.
17. Krizhevsky, A., Sutskever, I. & Hinton, G. E. ImageNet Classification with Deep Convolutional Neural Networks. In Pereira, F., Burges, C. J. C., Bottou, L. & Weinberger, K. Q. (eds.) *Advances in Neural Information Processing Systems 25*, 1097–1105 (Curran Associates, Inc., 2012).
18. Szegedy, C., Ioffe, S., Vanhoucke, V. & Alemi, A. A. Inception-v4, inception-resnet and the impact of residual connections on learning. In *AAAI*, vol. 4, 12 (2017).
19. Zhao, H., Zarar, S., Tashev, I. & Lee, C.-H. Convolutional recurrent neural networks for speech enhancement. *arXiv preprint arXiv:1805.00579* (2018).
20. Silver, D. *et al.* Mastering the game of go with deep neural networks and tree search. *Nature* **529**, 484–489 (2016).
21. Richardson, W. H. Bayesian-based iterative method of image restoration. *JOSA* **62**, 55–59 (1972).
22. McCann, M. T., Jin, K. H. & Unser, M. A review of convolutional neural networks for inverse problems in imaging. *arXiv preprint arXiv:1710.04011* (2017).
23. Lucas, A., Iliadis, M., Molina, R. & Katsaggelos, A. K. Using deep neural networks for inverse problems in imaging: beyond analytical methods. *IEEE Signal Processing Magazine* **35**, 20–36 (2018).
24. Hornik, K., Stinchcombe, M. & White, H. Multilayer feed-forward networks are universal approximators. *Neural networks* **2**, 359–366 (1989).
25. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986).
26. Masci, J., Meier, U., Cireşan, D. & Schmidhuber, J. Stacked convolutional auto-encoders for hierarchical feature extraction. In *International Conference on Artificial Neural Networks*, 52–59 (Springer, 2011).
27. Ronneberger, O., Fischer, P. & Brox, T. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 234–241 (Springer, 2015).
28. He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, 770–778 (IEEE Computer Society, 2016).
29. Goodfellow, I. *et al.* Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680 (2014).
30. Isola, P., Zhu, J., Zhou, T. & Efros, A. A. Image-to-image translation with conditional adversarial networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, 5967–5976 (IEEE Computer Society, 2017).
31. Carlton, P. M. *et al.* Fast live simultaneous multiwavelength four-dimensional optical microscopy. *Proceedings of the National Academy of Sciences* **107**, 16016–16022 (2010).
32. Marim, M. M., Angelini, E. D. & Olivo-Marin, J.-C. A compressed sensing approach for biological microscopy image denoising. In *SPARS'09-Signal Processing with Adaptive Sparse Structured Representations* (2009).
33. Boulanger, J. *et al.* Patch-based nonlocal functional for denoising fluorescence microscopy image sequences. *IEEE transactions on medical imaging* **29**, 442–454 (2010).
34. Luisier, F., Blu, T. & Unser, M. Image denoising in mixed poisson–gaussian noise. *IEEE Transactions on image processing* **20**, 696–708 (2011).
35. Xu, J., Zhang, L. & Zhang, D. A trilateral weighted sparse coding scheme for real-world image denoising. *arXiv preprint arXiv:1807.04364* (2018).
36. Yair, N. & Michaeli, T. Multi-scale weighted nuclear norm image restoration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3165–3174 (2018).

37. Buades, A., Coll, B. & Morel, J.-M. A non-local algorithm for image denoising. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 52, 113–147 (IEEE, 2005).

38. Dabov, K., Foi, A., Katkovnik, V. & Egiazarian, K. Image denoising with block-matching and 3d filtering. In *Image Processing: Algorithms and Systems, Neural Networks, and Machine Learning*, vol. 6064, 606414 (International Society for Optics and Photonics, 2006).

39. Jain, V. & Seung, S. Natural image denoising with convolutional networks. In *Advances in Neural Information Processing Systems*, 769–776 (2009).

40. Zhang, K., Zuo, W., Chen, Y., Meng, D. & Zhang, L. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing* **26**, 3142–3155 (2017).

41. Lehtinen, J. *et al.* Noise2noise: Learning image restoration without clean data. In Dy, J. G. & Krause, A. (eds.) *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, vol. 80 of *JMLR Workshop and Conference Proceedings*, 2971–2980 (JMLR.org, 2018).

42. Buchholz, T., Jordan, M., Pigino, G. & Jug, F. Cryo-care: Content-aware image restoration for cryo-transmission electron microscopy data. *CoRR* **abs/1810.05420** (2018). URL <http://arxiv.org/abs/1810.05420>. 1810.05420.

43. Batson, J. & Royer, L. Noise2self: Blind denoising by self-supervision. *arXiv preprint arXiv:1901.11365* (2019).

44. Krull, A., Buchholz, T. & Jug, F. Noise2void - learning denoising from single noisy images. *CoRR* **abs/1811.10980** (2018). URL <http://arxiv.org/abs/1811.10980>. 1811.10980.

45. Laine, S., Lehtinen, J. & Aila, T. Self-supervised deep image denoising. *arXiv preprint arXiv:1901.10277* (2019).

46. Ulyanov, D., Vedaldi, A. & Lempitsky, V. Deep image prior. *arXiv preprint arXiv:1711.10925* (2017).

47. Zhu, J., Park, T., Isola, P. & Efros, A. A. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, 2242–2251 (IEEE Computer Society, 2017).

48. Hom, E. F. *et al.* Aida: an adaptive image deconvolution algorithm with application to multi-frame and three-dimensional data. *JOSA A* **24**, 1580–1600 (2007).

49. Dey, N. *et al.* Richardson–lucy algorithm with total variation regularization for 3d confocal microscope deconvolution. *Microscopy research and technique* **69**, 260–266 (2006).

50. Preibisch, S. *et al.* Efficient bayesian-based multiview deconvolution. *Nature methods* **11**, 645–648 (2014).

51. Rivenson, Y. *et al.* Deep learning microscopy: Enhancing resolution, field-of-view and depth-of-field of optical microscopy images using neural networks. In *Conference on Lasers and Electro-Optics* (OSA, 2018).

52. Betzig, E. *et al.* Imaging intracellular fluorescent proteins at nanometer resolution. *Science (New York, N.Y.)* **313**, 1642–1645 (2006).

53. Rust, M. J., Bates, M. & Zhuang, X. Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (storm). *Nature methods* **3**, 793–795 (2006).

54. Henriques, R. *et al.* Quickpalm: 3d real-time photoactivation nanoscopy image processing in imagej. *Nature methods* **7**, 339–340 (2010).

55. Boyd, N., Jonas, E., Babcock, H. P. & Recht, B. Deeploco: Fast 3d localization microscopy using neural networks. *BioRxiv* 267096 (2018).

56. Gustafsson, N. *et al.* Fast live-cell conventional fluorophore nanoscopy with imagej through super-resolution radial fluctuations. *Nature communications* **7**, 12471 (2016).

57. Gustafsson, M. G. L. *et al.* Three-dimensional resolution doubling in wide-field fluorescence microscopy by structured illumination. *Biophysical journal* **94**, 4957–4970 (2008).

58. Li, D. *et al.* Extended-resolution structured illumination imaging of endocytic and cytoskeletal dynamics. *Science* **349**, aab3500 (2015).

59. Mudry, E. *et al.* Structured illumination microscopy using unknown speckle patterns. *Nat. Photonics* **6**, 312 (2012).

60. Ayuk, R. *et al.* Structured illumination fluorescence microscopy with distorted excitations using a filtered blind-SIM algorithm. *Opt. Lett.* **38**, 4723–4726 (2013).

61. Jahr, W., Schmid, B., Schmied, C., Fahrbach, F. O. & Huisken, J. Hyperspectral light sheet microscopy. *Nature communications* **6**, 7990 (2015).

62. Cutrale, F. *et al.* Hyperspectral phasor analysis enables multiplexed 5d in vivo imaging. *Nature methods* **14**, 149 (2017).

63. Blasse, C. *et al.* Premosa: extracting 2d surfaces from 3d microscopy mosaics. *Bioinformatics (Oxford, England)* **33**, 2563–2569 (2017).

64. Mayer, J., Robert-Moreno, A., Sharpe, J. & Swoger, J. Attenuation artifacts in light sheet fluorescence microscopy corrected by optispim. *Light, science & applications* **7**, 70 (2018).

65. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T. & Efros, A. A. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2536–2544 (2016).

66. Liu, G. *et al.* Image inpainting for irregular holes using partial convolutions <http://arxiv.org/abs/1804.07723v1>.

67. Amat, F. *et al.* Efficient processing and analysis of large-scale light-sheet microscopy data. *Nature Protocols* **10**, 1679–1696 (2015).

68. Cai, B., Xu, X., Jia, K., Qing, C. & Tao, D. Dehazenet: An end-to-end system for single image haze removal. *IEEE Transactions on Image Processing* **25**, 5187–5198 (2016).

69. Saalfeld, S., Fetter, R., Cardona, A. & Tomancak, P. Elastic volume reconstruction from series of ultra-thin microscopy sections. *Nature methods* **9**, 717–720 (2012).

70. Zbontar, J. & LeCun, Y. Computing the stereo matching cost with a convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1592–1599 (2015).

71. Rohé, M.-M., Datar, M., Heimann, T., Sermesant, M. & Pennec, X. Svf-net: Learning deformable image registration using shape matching. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 266–274 (Springer, 2017).

72. Nguyen, T., Chen, S. W., Skandan, S., Taylor, C. J. & Kumar, V. Unsupervised deep homography: A fast and robust homography estimation model. *IEEE Robotics and Automation Letters* (2018).

73. Christensen, R. P. *et al.* Untwisting the caenorhabditis elegans embryo. *eLife* **4** (2015).

74. Prevedel, R. *et al.* Simultaneous whole-animal 3d imaging of neuronal activity using light-field microscopy. *Nature methods* **11**, 727–730 (2014).

75. Fei, P. *et al.* Deep learning light field microscopy for rapid four-dimensional imaging of behaving animals. *bioRxiv* (2018).

76. Antipa, N. *et al.* Diffusercam: lensless single-exposure 3d imaging. *Optica* **5**, 1–9 (2018).

77. Vinegoni, C., Pitsouli, C., Razansky, D., Perrimon, N. & Ntziachristos, V. In vivo imaging of drosophila melanogaster pupae with mesoscopic fluorescence tomography. *Nature methods* **5**, 45–47 (2008).

78. Xingjian, S. *et al.* Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*, 802–810 (2015).

79. Naganathan, S. R., Fürthauer, S., Nishikawa, M., Jülicher, F. & Grill, S. W. Active torque generation by the actomyosin cell cortex drives left-right symmetry breaking. *eLife* **3**, e04165 (2014).

80. Meister, S., Hur, J. & Roth, S. Unflow: Unsupervised learning of optical flow with a bidirectional census loss. *arXiv preprint arXiv:1711.07837* (2017).

81. Haring, M. T. *et al.* Automated sub-5 nm image registration in integrated correlative fluorescence and electron microscopy using cathodoluminescence pointers. *Scientific reports* **7**, 43621 (2017).

82. Royer, L. A. *et al.* Adaptive light-sheet microscopy for long-term, high-resolution imaging in living organisms. *Nat. Biotechnol.* **34**, 1267–1278 (2016).

83. Liu, T.-L. *et al.* Observing the cell in its native state: Imaging subcellular dynamics in multicellular organisms. *Science (New York, N.Y.)* **360** (2018).

84. Turpin, A., Vishniakou, I. & Seelig, J. D. Light scattering control with neural networks in transmission and reflection. *arXiv preprint arXiv:1805.05602* (2018).

85. Recht, B., Roelofs, R., Schmidt, L. & Shankar, V. Do CIFAR-10 classifiers generalize to cifar-10? *CoRR* **abs/1806.00451** (2018). URL <http://arxiv.org/abs/1806.00451>.

86. Sabour, S., Cao, Y., Faghri, F. & Fleet, D. J. Adversarial manipulation of deep representations. *CoRR* **abs/1511.05122** (2015). URL <http://arxiv.org/abs/1511.05122>.

87. Osokin, A., Chessel, A., Salas, R. E. C. & Vaggi, F. Gans for biological image synthesis. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, 2252–2261 (IEEE, 2017).

88. Yosinski, J., Clune, J., Bengio, Y. & Lipson, H. How transferable are features in deep neural networks? In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D. & Weinberger, K. Q. (eds.) *Advances in Neural Information Processing Systems* **27**, 3320–3328 (Curran Associates, Inc., 2014).

89. Zeiler, M. D. & Fergus, R. Visualizing and Understanding Convolutional Networks. In *Computer Vision – ECCV 2014*, 818–833 (Springer International Publishing, 2014).

90. Esteva, A. *et al.* Dermatologist-level classification of skin cancer with deep neural networks. *Nature* **542**, 115 (2017).

91. Thul, P. J. *et al.* A subcellular map of the human proteome. *Science (New York, N.Y.)* **356** (2017).

92. Ljosa, V., Sokolnicki, K. L. & Carpenter, A. E. Annotated high-throughput microscopy image sets for validation. *Nature methods* **9**, 637 (2012).

93. Sullivan, D. P. *et al.* Deep learning is combined with massive-scale citizen science to improve large-scale image classification. *Nature biotechnology* **36**, 820 (2018).
94. Simonyan, K., Vedaldi, A. & Zisserman, A. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034* (2013).
95. Zhang, Q., Cao, R., Shi, F., Wu, Y. N. & Zhu, S.-C. Interpreting cnn knowledge via an explanatory graph. *arXiv preprint arXiv:1708.01785* (2017).
96. Zhang, Q., Wu, Y. N. & Zhu, S.-C. Interpretable convolutional neural networks. *arXiv preprint arXiv:1710.00935* **2**, 5 (2017).
97. Lakshminarayanan, B., Pritzel, A. & Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, 6402–6413 (2017).
98. Hutson, M. Artificial intelligence faces reproducibility crisis. *Science (New York, N.Y.)* **359**, 725–726 (2018).
99. Henderson, P. *et al.* Deep reinforcement learning that matters. *arXiv preprint arXiv:1709.06560* (2017).

**Acknowledgements** We would like to thank our colleagues at the CZ Biohub: Anitha Krishnan, Bryant Chhun, Manuel Leonetti, Shalin Mehta, Josh Batson, and Rafael Gomez Sjoberg for insightful discussions, feedback and for reviewing the manuscript. Thanks to our colleague James Zou from Stanford University for advice. We would like to thank Martin Weigert for reviewing the manuscript and for innumerable discussions on the topic of Deep Learning and microscopy. Thanks to Wei Ouyang, Robert Prevedel, and Florian Jug for feedback on the first version of the preprint. Funding for this work was provided by the Chan Zuckerberg Biohub.

**Competing Interests** The authors declare that they have no competing financial interests.

**Correspondence** Correspondence and requests for materials should be addressed to Loic A. Royer (email: loic.royer@czbiohub.org, twitter: @loicaroyer).