

Artificial Intelligence for Microscopy: What You Should Know

Lucas von Chamier¹, Romain F. Laine^{1-3, ✉}, and Ricardo Henriques^{1-3,4 ✉}

¹MRC-Laboratory for Molecular Cell Biology, University College London, London, UK

²Department of Cell and Developmental Biology, University College London, London, UK

³The Francis Crick Institute, London, UK

⁴Institute for the Physics of Living Systems, University College London, London, UK

Artificial Intelligence based on Deep Learning is opening new horizons in biomedical research and promises to revolutionize the microscopy field. Slowly, it now transitions from the hands of experts in computer sciences to researchers in cell biology. Here, we introduce recent developments in deep learning applied to microscopy, in a manner accessible to non-experts. We overview its concepts, capabilities and limitations, presenting applications in image segmentation, classification and restoration. We discuss how deep learning shows an outstanding potential to push the limits of microscopy, enhancing resolution, signal and information content in acquired data. Its pitfalls are carefully discussed, as well as the future directions expected in this field.

Artificial intelligence | Machine learning | Live-cell imaging | Super-resolution microscopy | Classification | Segmentation


Correspondence: r.laine@ucl.ac.uk, r.henriques@ucl.ac.uk

Introduction. Deep Learning is a type of Artificial Intelligence (AI) which has recently seen a rise in academic research and popular interest. This sudden boost has been primarily fuelled by the invention of Convolutional Neural Networks (CNNs), a novel machine learning algorithmic architecture. In the early 2010's CNNs became increasingly prominent as tools for image classification, showing super-human accuracy at identifying objects in images (1). Since then, Deep Learning has expanded to many research fields, showing its potential to outsmart humans in board games such as Go (2), achieve self-driving cars (3, 4) and to significantly improve biomedical image analysis (5).

Recently, a growing amount of biology studies has used CNNs to analyze microscopy data, laying the foundation for a fundamental change in how imaging data is interpreted, and how microscopy is carried out. The areas where AI has been applied include: automated, accurate classification and segmentation of microscopy images (6–8); extraction of structures from label-free microscopy imaging (artificial labelling) (9, 10); and image restoration, such as denoising and resolution enhancement (11–13). Here, we give non-specialist readers an overview of the potential of Deep Learning, specifically through CNNs, in the context of some of the major challenges of microscopy. We will also discuss some of the current limitations of the approach and give an outlook on possible future applications of Deep Learning in microscopy.

For a more in-depth view into the AI-field, we advise the reader to see the review by Lecun *et al.* (14), which gives an extended perspective on Deep Learning and its historical de-

a. Classical algorithm



b. Deep Learning




Fig. 1. Deep Learning compared to classical computation. a) Classical computer programs convert an input (e.g. blurry image) into a desired output (sharp image) via an algorithm with known rules and parameters. b-i) Neural networks are trained with pairs of known inputs and outputs, e.g. a blurry and highly resolved image of a cell. During training, the network learns to match its inputs by observing a large number of paired examples. While doing so, it optimizes its internal parameters, gaining the capacity to predict a matching image when given a single input (b-ii). b-ii) After training, the network can be used to perform the task similarly to a conventional algorithm on novel data.

velopment. Additionally Litjens *et al.* (15), Angermueller *et al.* (16) and Belthangady *et al.* (17) comprehensively discuss the application of AI in Biomedical Sciences and Computational Biology.

How does a CNN learn? CNNs are complex networks of connected ‘neurons’ arranged in ‘layers’, nomenclature inspired by the animal visual cortex (18). Each neuron performs specific mathematical operations on its input and passes the output to the neurons of the next layer of the network. In CNNs, the operations are typically either convolutions which extract patterns from images, or so-called

pooling operations which reduce the number of pixels in the image and therefore simplify the data representations. This combination of feature extractions and data simplification is what allows CNNs to 'understand' the content of the images and perform efficiently on many types of imaging data. Once several layers are stacked upon another, networks become 'deep' and the information they can extract from inputs becomes increasingly complex (14).

Unlike conventional computer programs which are designed to perform a set of user-defined and well-understood operations on the input data (Fig. 1a) (19), a CNN has to be trained (Fig. 1b-i) on a so-called training dataset. This set consists of paired inputs and outputs which describe the transformation of the image that the user wishes the network to learn. For instance, for an image denoising operation, noisy images and matching high signal images have to be provided for training. The training dataset therefore defines the task that the network performs. During training, the network compares its output to the provided input and adjusts the parameters of its neurons, e.g. the impact that each neuron has on the next layer, until the network is able to infer an artificial output which resembles the expected real output. After training, networks are tested on a validation dataset, using inputs previously 'unseen' by the network, to determine if it can generalize over new data. In this case, real outputs are not used to train the network but to determine its performance. Once trained and validated, the network can be applied to new data for which no real output exists, identically to conventional computer programs (Fig. 1b-ii). Training is computationally intensive, taking hours to days, especially for deep networks with millions of trainable parameters. In comparison, after training, the inference process is considerably fast, taking minutes to seconds.

However, the computational performance of CNNs is incrementally improving with the development of increasingly powerful processing units, notably Graphical Processing Units (GPUs). Until the introduction of the first GPU-enabled CNN in 2012 (1) called AlexNet, CNNs were largely neglected in AI (14), because their training was too slow, sometimes requiring weeks to months of computation to complete. AlexNet greatly outperformed the competition at the ImageNet image classification challenge in 2012, a seminal breakthrough for the AI field.

The success of CNN algorithms depends on the design of the network architecture, usually carried out by computer scientists, and the availability of good training data. Generally, the training dataset should contain many different examples of the desired outputs. For example, a network designed to learn to categorize an animal should be trained with images showing the animal in different positions or environments. Generating and curating the training dataset is often the major bottleneck for the application of CNNs.

CNNs in microscopy. One of the first studies using GPU-enabled neural networks was able to segment neuronal membranes from electron microscopy (EM) images (20), vastly improving speed and efficiency of segmentation over the state-of-art methods. Another breakthrough came through

the design of a more efficient network architecture which combines a number of convolution/pooling layers (the encoder), with a number of layers of de-convolution/up-sampling (the decoder) (7, 21). The encoder learns the main features of the image and the decoder reassigns them to different pixels of the image. Due to this conceptual down- followed by up-sampling, this network architecture was termed 'U-net'. U-nets are therefore very powerful for image-to-image tasks (as opposed to simple classification of the image), making them one of the most important networks for microscopy applications (5, 9, 11, 22).

Researchers in life sciences face several challenges when imaging biological specimens: How can phototoxicity and bleaching of fluorescent labels be balanced against good signal or resolution? How many channels can a cell be imaged in without interfering with native processes? And how can relevant and complex information be extracted from large image datasets, without tedious manual annotation and human bias? In the following sections, we will present how AI methods have recently provided efficient solutions to these problems. While there exist some conceptual overlaps, we have separated these into four categories: image classification, image segmentation, artificial labelling and image restoration.

Classification. An important goal for microscopic image analysis is to recognize and assign identities to relevant features on an image (Fig. 2). For example, identifying mitotic cells in a tissue sample can be essential for cancer diagnosis. However, manual annotation is tedious, limited in throughput, and experts can introduce bias into such annotations by deciding which image features are important while ignoring others. Although several computational methods have been introduced to accelerate detection or classification tasks (23–25), these still often rely on handcrafted parameters, chosen by researchers. The advantage of CNNs is their capability to learn relevant image features automatically. CNNs have therefore been extensively used in the Biomedical imaging field, especially for cancer detection, particularly as large training sets have become more available (6, 26–30). The classification accuracy of neural networks has been successfully applied to high-throughput and high-content screens where it has shown expert-level recognition of subcellular features (31–34). Producing hand-labelled classification examples for such tasks can be a constraint for the creation of training sets. However, Lu et al. have shown that this can be overcome with unsupervised learning, where a network




Fig. 2. Classification Schematic of neural network trained to detect and classify cells of different types or stages, e.g. to identify mitotic cells

learns to predict protein localisations by being shown paired

cells which were either imaged or not imaged in a specific channels, requiring no manual labelling to produce training inputs. (35).

CNNs have also shown the capacity to accurately identify cellular states from transmitted-light data, for example, differentiating cells based on cell cycle stage (36), cells affected by phototoxicity (37) or stem cell-derived endothelial cells (38). Identifying such identities, particularly in large image-sets previously often required the introduction of a label into the cell, an obstacle which CNNs can now overcome.

Segmentation. Segmentation is the identification of image regions that are part of specific cellular or sub-cellular structures and often is an essential step in image analysis (Fig. 3). A drawback of some existing segmentation platforms (23, 25, 39) is that they need user-based fine-tuning and manual error-removal, requiring time and expertise (40). Image segmentation can therefore be a bottleneck for research, particularly for high-throughput studies. CNNs have in multiple studies outperformed classical approaches in terms of accuracy and generalization (7, 20, 21, 40, 41), especially when performing cell segmentation in co-cultures of multiple cell types (40). In the context of histopathology, CNNs have been successfully used to segment colon glands (42–46) and breast tissues (47, 48), outperforming non-deep learning based approaches.

Naturally, there is overlap between the challenges of classification and segmentation since both require the network to learn about specific regions of an image. Hence, segmentation is often used with subsequent classification and can even improve the accuracy of classification (28, 49).




Fig. 3. Segmentation Schematic of neural network trained to assign color-values to pixels on images, allowing the construction of segmentation masks.

Artificial Labelling. The analysis of specific structures in cells, especially in light microscopy, typically requires the introduction of labels, either by genetic labelling or chemical staining, which can disturb the biological system. Additionally, fluorescence microscopy, especially when using laser illumination, is inherently more phototoxic to cells than transmitted-light imaging (50, 51). With this in mind, two studies using CNN methods have shown that specific cellular structures, such as nuclear membrane, nucleoli, plasma membranes and mitochondria, can be extracted by neural networks from label-free images (9, 10). The networks used here were trained to predict a fluorescent label from transmitted-light or EM images, alleviating the need to label and acquire the corresponding fluorescence images (Fig. 4). This capacity is especially useful when tracking individual structures

over long periods of time. It was also shown that the networks can achieve high accuracy using a training dataset of only 30 to 40 images (9), and can simultaneously identify dying cells or distinguish different cell-types and subcellular structures (10). Christiansen *et al.* (10) also demonstrated their network's ability for transfer learning, a method allowing a pre-trained network to perform a new task with minimal additional training. The network showed promising results for transfer learning between different microscopes and labels, highlighting the versatility of these networks' performance. While the task of artificial labelling is similar to 'classic' segmentation, the main difference in this approach lies in the creation of the training set which does not require to be hand-labelled. The networks are simply trained from paired frames obtained from cells imaged in bright-field and fluorescence modalities.




Fig. 4. Artificial Labelling Schematic of neural network trained to label cellular structures in images of unstained cells or cells imaged in bright field.

Image Restoration: resolution and signal. The amount and quality of features which can be extracted from a microscopic image is limited by fundamental constraints inherent to all optical set-ups: signal-to-noise ratio (SNR) and resolution. Therefore, overcoming these limitations constitutes a central research goal in microscopy. In particular, super-resolution microscopy (SRM) (52–56) now allows imaging of cellular structures at the nanoscale using light microscopy. However, phototoxicity, bleaching and low temporal resolution still limit the capacity to achieve high-resolution long-term imaging in living specimens. Several research groups have proposed CNN methods addressing some of these issues in the last two years (11, 13?). Training approaches for networks published to date fall into two categories which could be described as naïve and informed. In the naïve approach, training datasets consist, for instance, of pairs of images acquired at low and high SNR and the network learns to predict a denoised image from a noisy input. The training dataset can be obtained from acquiring images at short and long exposure time or at different laser intensities. This approach was demonstrated on the highly photosensitive organism *Schmidtea mediterranea* and allowed a 60-fold decrease in illumination dose, enabling longer and more detailed observation of this organism *in vivo* (11). The authors also used the same approach to obtain SRM images from diffraction-limited image data, using the SRRF method as a reference (56, 57). Similarly, SRM images can be obtained from conventional confocal microscopy images using STED to acquire the high-resolution training dataset (?).

CNNs in Single-Molecule Localization Microscopy.

CNNs have also recently generated interest in the Single-molecule Localization Microscopy (SMLM) field. All available studies were published within the last year, by independent groups, suggesting that the potential of AI for SRM is increasingly recognized in the community (13, 58–61). Applying sophisticated network architectures, with combinations of widefield and SMLM data as inputs (13, 58), it is possible to train a CNN for SMLM reconstruction. Here, the networks do not learn to localize individual fluorophores as typical single-molecule localization algorithms do, but instead to map sparse SMLM data of either microtubules, mitochondria or nuclear pores into SRM output images. This demonstrates the strength of CNNs for pattern recognition in redundant data, like SMLM data where only a small number of frames may suffice to reconstruct a Super-Resolution image. Interestingly, some of these algorithms require no parameter tuning or specific knowledge about the imaged structures (58). Especially, for high emitter density, this is advantageous over conventional SMLM reconstruction algorithms which can be time-consuming and require sample-dependent optimisation of imaging parameters.

Two other groups have used a different approach to SMLM reconstruction which could be termed ‘informed’ as it makes use of the intrinsic properties of SMLM data (59, 60). Here, networks are trained to learn the positions of fluorophores from SMLM input images. In contrast to naïve training, this means that such CNNs can only be used for data acquired with SMLM techniques. However, the advantage is that the output of these networks contains the positional information of the fluorophores whereas the output of naïve networks does not. The outstanding question about how neural networks manage to produce Super Resolution images from sparse or widefield data, is circumvented by this approach. The reconstructed images are therefore more similar to standard SMLM reconstructions making the resolution improvement easier to interpret.

While achieving similar accuracy to state of the art SMLM algorithms (62), a main achievement of deep learning for SMLM is the speed with which super-resolved images can be produced. In several studies this was increased by several orders of magnitude compared to conventional reconstruction algorithms (13, 58, 59, 61).

Discussion. The use of neural networks is transforming microscopy both by allowing human or super-human performances for a number of image analysis tasks and as an automated high-performance tool for big-data analysis (33, 34, 40) (Table 1). However, while performance, versatility and speed of neural networks is likely to continue increasing, there are significant challenges which will not be solved by improved processing units. A frequently raised concern in the microscopy community over AI is how much machine outputs can be trusted to represent data. This is a real concern since CNNs have been observed to cause image hallucinations (67) or to fail catastrophically simply as a result of minute changes in the image (68). To address this issue, several groups have assessed the presence of artefacts in their

their network output images, notably using SQUIRREL (11–13, 61, 69, 70).

While this may identify the presence of artefacts, it does not address the underlying problem that it is difficult to interpret how CNN architectures produce their output from the image input. This lack of interpretability of network outputs is particularly concerning in the case of resolution enhancement, where it is not clear what information a CNN can extract from a diffraction limited image to achieve a non-diffraction limited image and how deep learning algorithms achieve this without producing significantly more artefacts than standard algorithms (11, 12). Another consequence of this is that the design of CNN architectures has been referred to as ‘notorious as an empirical endeavour’ (10). Choosing network parameters such as network depth, number of neural connections, learning rate and other hand-coded features of neural networks, also termed hyperparameters, is therefore based on evaluating the respective choices by performance (33, 60, 71, 72).

Beside issues of interpretability, there are other anecdotal examples where networks have ‘cheated’ their way to high performance, e.g. by using undesirable features such as empty space (37) to identify dead cells or by identifying patterns in the ordering of the training set, but not in the images themselves (73). This shows how much of the performance of Deep Learning methods relies on the choice and curation of training data sets, which can be a significant challenge for the application of deep learning for microscopy.

Despite these issues, AI has great enabling potential for microscopy, given super-human performance in classification tasks and image reconstruction. Hence, the issues discussed above should not discourage the use of CNNs as a research tool but be reason for caution when interpreting the performance of neural networks, as for any computational analysis tool.

Outlook Most of the methods reviewed here were published within the past 24 months, demonstrating the massive interest in deep learning as a versatile and powerful tool for microscopy applications. However, the delay between developments and their applications means that some areas of AI research have not yet been widely translated to microscopy. For example, transfer learning is an area which will likely become more widely investigated, allowing the use of pre-trained networks to carry out a new task, forms of which are only starting to become available (8). Finding methods to reuse neural networks robustly on multiple different tasks, different image sizes or images taken on different microscopes would make deep learning a more flexible and usable approach for image analysis than is currently possible. Importantly, it would reduce the need for large training datasets and shorten the training needed for new tasks. It could therefore lower the accessibility barrier of the approach and minimize the need for biologists to be fully familiar with neural network specifics. In turn, this would allow deep learning to become a tool, rather than a method that needs expert knowledge. This will likely open the door for deep learning to become even more widely used within life science.

Table 1. Major publications from the recent years applying or developing deep learning to microscopy. This table covers the main four themes where AI has provided solutions to some of the major limitations of microscopy in the recent years.

Classification/Segmentation			
Authors	Year	Details	Reference
Ning <i>et al.</i>	2005	Classification and segmentation of tissues in stages of <i>C. elegans</i> development	(41)
Ciresan <i>et al.</i>	2012	Segmentation of neuronal membranes in EM	(20)
Ciresan <i>et al.</i>	2013	Mitosis detection in breast cancer	(6)
Long <i>et al.</i>	2015	Introduction of fully convolutional networks (fCNN) for segmentation tasks	(21)
Ronneberger <i>et al.</i>	2015	U-net: significant increase in efficiency for image to image networks	(7)
Kainz <i>et al.</i>	2015	Colon gland segmentation, network outperforms handcrafted feature detection	(42)
Kraus <i>et al.</i>	2016	High-throughput classification and segmentation in yeast	(32)
Z. Xu <i>et al.</i>	2016	Accelerated cell detection in very large images and images with large cell numbers	(26)
Bauer <i>et al.</i>	2016	Cancer cell survival in different organs	(63)
Yao <i>et al.</i>	2016	Lung cancer survival prediction via biomarker (subtype) discovery	(64)
Dürr <i>et al.</i>	2016	Phenotyping of over 40000 drug-treated single-cell images	(65)
van Valen <i>et al.</i>	2016	Efficient cell-cell segmentation in live-cell imaging	(40)
Richmond <i>et al.</i>	2017	Identification of cells showing damage from phototoxicity	(37)
Kraus <i>et al.</i>	2017	Identification of yeast strains and mutants and subcellular protein localisations	(34)
M. Xu <i>et al.</i>	2017	High-throughput identification of blood cells from sickle-cell anemia samples	(66)
Pärnamaa & Parts	2017	High-throughput classification of protein localisation in yeast	(33)
Godinez <i>et al.</i>	2017	Phenotyping cells after drug treatment, organelle identification	(31)
Eulenberg <i>et al.</i>	2017	Identification of cell-cycle phases and differences in disease stages	(36)
Kusumoto <i>et al.</i>	2018	Detection of epithelial cells derived from iPSCs	(38)
Lu <i>et al.</i>	2018	Unsupervised learning for protein localisation prediction, in human and yeast cells	(35)
Falk <i>et al.</i>	2018	U-net as ImageJ plugin for non-experts, includes pretrained network	(8)

Artificial Labelling			
Authors	Year	Details	Reference
Christiansen <i>et al.</i>	2018	Label Prediction in fixed and live cells	(10)
Ounkomol <i>et al.</i>	2018	3D label prediction in live-cell, IF and EM images	(9)

Image Restoration/Super Resolution			
Authors	Year	Details	Reference
Nehme <i>et al.</i>	2018	SMLM images from diffraction limited input	(58)
Boyd <i>et al.</i>	2018	Identifies localisation of fluorophores from STORM single frames	(59)
Ouyang <i>et al.</i>	2018	SMLM reconstruction using very small number of frames to predict SR image	(13)
Nelson & Hess	2018	PALM reconstruction network trained directly on the image to be analysed	(60)
Weigert <i>et al.</i>	2018	Denoising and resolution enhancement in different organisms and cell types	(11)
Li <i>et al.</i>	2018	SMLM localisation by deep learning and artefact removal by statistical inference	(61)
Wang <i>et al.</i>	2019	Conversion of low NA to high NA or diffraction limited to STED resolution images	(12)

ACKNOWLEDGEMENTS

We'd like to thank Simon F. Nørrelykke (ETH Zürich, Switzerland) and Alex Lu (University of Toronto, Canada) for kindly suggesting corrections to this manuscript, see [here](#). This work was funded by grants from the UK Biotechnology and Biological Sciences Research Council (BB/M022374/1; BB/P027431/1; BB/R000697/1; BB/S507532/1) (R.H. and R.F.L.), the UK Medical Research Council (MR/K015826/1) (R.H.), the Wellcome Trust (203276/Z/16/Z) (R.H.), Core funding to the MRC Laboratory for Molecular Cell Biology, University College London (MC_UU12018/7). L.C supported by a 4-year MRC Research Studentship.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

Bibliography

- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354, 2017.
- Ana I Maqueda, Antonio Loquercio, Guillermo Gallego, Narciso García, and Davide Scaramuzza. Event-based vision meets deep learning on steering prediction for self-driving cars. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5419–5427, 2018.
- Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Praseon Goyal, Lawrence D Jackel, Mathew Monfort, Urs Müller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- Holger R Roth, Chen Shen, Hirohisa Oda, Masahiro Oda, Yuichiro Hayashi, Kazunari Misawa, and Kensaku Mori. Deep learning and its application to medical image segmentation. *Medical Imaging Technology*, 36(2):63–71, 2018.
- Dan C Cireşan, Alessandro Giusti, Luca M Gambardella, and Jürgen Schmidhuber. Mitosis detection in breast cancer histology images with deep neural networks. In *International Conference on Medical Image Computing and Computer-assisted Intervention*, pages 411–418. Springer, 2013.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- Thorsten Falk, Dominic Mai, Robert Bensch, Özgün Çiçek, Ahmed Abdulkadir, Yassine Marrakchi, Anton Böhmer, Jan Deubner, Zoe Jäckel, Katharina Seiwald, et al. U-net: deep

- learning for cell counting, detection, and morphometry. *Nature methods*, page 1, 2018.
9. Chawin Ounkomol, Sharmishta Seshamani, Mary M Maleckar, Forrest Collman, and Gregory R Johnson. Label-free prediction of three-dimensional fluorescence images from transmitted-light microscopy. *Nature Methods*, 15(11):917–920, 2018.
 10. Eric M Christiansen, Samuel J Yang, D Michael Ando, Ashkan Javaherian, Gaia Skibinski, Scott Lipnick, Elliot Mount, Alison O’Neil, Kevan Shah, Alicia K Lee, et al. In silico labeling: Predicting fluorescent labels in unlabeled images. *Cell*, 173(3):792–803, 2018.
 11. Martin Weigert, Uwe Schmidt, Tobias Boothe, Andreas Müller, Alexandr Dibrov, Akanksha Jain, Benjamin Wilhelm, Deborah Schmidt, Coleman Broaddus, Siân Culley, et al. Content-aware image restoration: pushing the limits of fluorescence microscopy. *Nature methods*, 15(12):1090, 2018.
 12. Hongda Wang, Yair Rivenson, Yiyin Jin, Zhensong Wei, Ronald Gao, Harun Günaydin, Laurent A Bentolila, Comert Kural, and Aydogan Ozcan. Deep learning enables cross-modality super-resolution in fluorescence microscopy. *Nature Methods*, 16(1):103–110, 2019.
 13. Wei Ouyang, Andrey Aristov, Mickaël Lelek, Xian Hao, and Christophe Zimmer. Deep learning massively accelerates super-resolution localization microscopy. *Nature biotechnology*, 2018.
 14. Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436, 2015.
 15. Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen Awm Van Der Laak, Bram Van Ginneken, and Clara I Sánchez. A survey on deep learning in medical image analysis. *Medical image analysis*, 42:60–88, 2017.
 16. Christof Angermueller, Tanel Pärnamaa, Leopold Parts, and Oliver Stegle. Deep learning for computational biology. *Molecular systems biology*, 12(7):878, 2016.
 17. Chinmay Belthangady and Loic A Royer. Applications, promises, and pitfalls of deep learning for fluorescence image reconstruction. 2018.
 18. F Rosenblatt. THE PERCEPTRON : A PROBABILISTIC MODEL FOR INFORMATION STORAGE AND ORGANIZATION. *Psychological Review*, 65(6):386–408, 1958.
 19. Ann Wheeler and Ricardo Henriques. Standard and super-resolution bioimaging data analysis: A primer, 2017.
 20. Dan Ciresan, Alessandro Giusti, Luca M Gambardella, and Jürgen Schmidhuber. Deep neural networks segment neuronal membranes in electron microscopy images. In *Advances in neural information processing systems*, pages 2843–2851, 2012.
 21. Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
 22. Jianxu Chen, Lin Yang, Yizhe Zhang, Mark Alber, and Danny Z Chen. Combining fully convolutional and recurrent neural networks for 3d biomedical image segmentation. In *Advances in Neural Information Processing Systems*, pages 3036–3044, 2016.
 23. Anne E Carpenter, Thouis R Jones, Michael R Lamprecht, Colin Clarke, In Han Kang, Ola Friman, David A Guertin, Joo Han Chang, Robert A Lindquist, Jason Moffat, Polina Golland, and David M Sabatini. CellProfiler : image analysis software for identifying and quantifying cell phenotypes. *Genome Biology*, 7(10), 2006. doi: 10.1186/gb-2006-7-10-r100.
 24. Michael Held, Michael H A Schmitz, Bernd Fischer, Thomas Walter, Beate Neumann, Michael H Olma, Matthias Peter, Jan Ellenberg, and Daniel W Gerlich. CellCognition : time-resolved phenotype annotation in high-throughput live cell imaging. *Nature Publishing Group*, 7(9):747–754, 2010. ISSN 1548-7091. doi: 10.1038/nmeth.1486.
 25. Christoph Sommer, Christoph N Straehle, Ullrich Koethe, Fred A Hamprecht, et al. Ilastik: Interactive learning and segmentation toolkit. In *ISBI*, volume 2, page 8, 2011.
 26. Zheng Xu and Junzhou Huang. Detecting 10,000 cells in one second. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 676–684. Springer, 2016.
 27. Christopher D Malon and Eric Cosatto. Classification of mitotic figures with convolutional neural networks and seeded blob features. *Journal of pathology informatics*, 4, 2013.
 28. Dayong Wang, Aditya Khosla, Rishab Gargeya, Humayun Irshad, and Andrew H Beck. Deep learning for identifying metastatic breast cancer. *arXiv preprint arXiv:1606.05718*, 2016.
 29. Anat Shkolyar, Amit Gefen, Dafna Benayahu, and Hayit Greenspan. Automatic detection of cell divisions (mitosis) in live-imaging microscopy images using convolutional neural networks. In *Engineering in Medicine and Biology Society (EMBC), 2015 37th Annual International Conference of the IEEE*, pages 743–746. IEEE, 2015.
 30. Yan Xu, Tao Mo, Qiwei Feng, Peilin Zhong, Maode Lai, and Eric I.Chao Chang. Deep learning of feature representation with multiple instance learning for medical image analysis. *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, (1):1626–1630, 2014. doi: 10.1109/ICASSP.2014.6853873.
 31. William J Godinez, Imtiaz Hossain, Stanley E Lazic, John W Davies, and Xian Zhang. A multi-scale convolutional neural network for phenotyping high-content cellular images. *Bioinformatics*, 33(13):2010–2019, 2017.
 32. Oren Z Kraus, Jimmy Lei Ba, and Brendan J Frey. Classifying and segmenting microscopy images with deep multiple instance learning. *Bioinformatics*, 32(12):i52–i59, 2016.
 33. Tanel Pärnamaa and Leopold Parts. Accurate classification of protein subcellular localization from high-throughput microscopy images using deep learning. *G3: Genes, Genomes, Genetics*, 7(5):1385–1392, 2017.
 34. Oren Z Kraus, Ben T Grys, Jimmy Ba, Yolanda Chong, Brendan J Frey, Charles Boone, and Brenda J Andrews. Automated analysis of high-content microscopy data with deep learning. *Molecular systems biology*, 13(4):924, 2017.
 35. Alex Lu, Oren Z Kraus, Sam Cooper, and Alan M Moses. Learning unsupervised feature representations for single cell microscopy images with paired cell inpainting. *bioRxiv*, page 395954, 2018.
 36. Philipp Eulenberg, Niklas Köhler, Thomas Blasi, Andrew Filby, Anne E Carpenter, Paul Rees, Fabian J Theis, and F Alexander Wolf. Reconstructing cell cycle and disease progression using deep learning. *Nature communications*, 8(1):463, 2017.
 37. David Richmond, Anna Payne-Tobin Jost, Talley Lambert, Jennifer Waters, and Hunter Elliott. Deadnet: identifying phototoxicity from label-free microscopy images of cells using deep convnets. *arXiv preprint arXiv:1701.06109*, 2017.
 38. Dai Kusumoto, Mark Lachmann, Takeshi Kunihiro, Shinsuke Yuasa, Yoshikazu Kishino, Mai Kimura, Toshiomi Katsuki, Shogo Itoh, Tomohisa Seki, and Keiichi Fukuda. Automated deep learning-based system to identify endothelial cells derived from induced pluripotent stem cells. *Stem cell reports*, 10(6):1687–1695, 2018.
 39. Ignacio Arganda-Carreras, Verena Kastnig, Curtis Rueden, Kevin W Eliceiri, Johannes Schindelin, Albert Cardona, and H Sebastian Seung. Trainable weka segmentation: a machine learning tool for microscopy pixel classification. *Bioinformatics*, 33(15):2424–2426, 2017.
 40. David A Van Valen, Takamasa Kudo, Keara M Lane, Derek N Macklin, Nicolas T Quach, Mialy M DeFelicis, Inbal Maayan, Yu Tanouchi, Euan A Ashley, and Markus W Covert. Deep learning automates the quantitative analysis of individual cells in live-cell imaging experiments. *PLoS computational biology*, 12(11):e1005177, 2016.
 41. Feng Ning, Damien Delhomme, Yann Lecun, Fabio Piano, Léon Bottou, Paolo Emilio Barbano, and A Automatic Phenotyping. Toward Automatic Phenotyping of Developing Embryos From Videos. *IEEE Transactions on Image Processing*, 14(9):1360–1371, 2005.
 42. Philipp Kainz, Michael Pfeiffer, and Martin Urschler. Semantic segmentation of colon glands with deep convolutional neural networks and total variation segmentation. *arXiv preprint arXiv:1511.06919*, 2015.
 43. Aicha BenTaieb and Ghassan Hamarneh. Topology aware fully convolutional networks for histology gland segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 460–468. Springer, 2016.
 44. Hao Chen, Xiaojun Qi, Lequan Yu, and Pheng-Ann Heng. Dcan: deep contour-aware networks for accurate gland segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2487–2496, 2016.
 45. Wenqi Li, Siyamalan Manivannan, Shazia Akbar, Jianguo Zhang, Emanuele Trucco, and Stephen J McKenna. Gland segmentation in colon histology images using hand-crafted features and convolutional neural networks. In *Biomedical Imaging (ISBI), 2016 IEEE 13th International Symposium on*, pages 1405–1408. IEEE, 2016.
 46. Yan Xu, Yang Li, Mingyuan Liu, Yipei Wang, Maode Lai, I Eric, and Chao Chang. Gland instance segmentation by deep multichannel side supervision. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 496–504. Springer, 2016.
 47. Jun Xu, Xiaofei Luo, Guan hao Wang, Hannah Gilmore, and Anant Madabhushi. A deep convolutional neural network for segmenting and classifying epithelial and stromal regions in histopathological images. *Neurocomputing*, 191:214–223, 2016.
 48. Geert Litjens, Clara I Sánchez, Nadya Timofeeva, Mayke Hermesen, Iris Nagtegaal, Iringo Kovacs, Christina Hulsbergen-Van De Kaa, Peter Bult, Bram Van Ginneken, and Jeroen Van Der Laak. Deep learning as a tool for increased accuracy and efficiency of histopathological diagnosis. *Scientific reports*, 6:26286, 2016.
 49. Claire Lian Chen, Ata Mahjoubfar, Li-Chia Tai, Ian K Blaby, Allen Huang, Kayvan Reza Niazi, and Bahram Jalali. Deep learning in label-free cell classification. *Scientific reports*, 6:21471, 2016.
 50. Ram Dixit and Richard Cyr. Cell damage and reactive oxygen species production induced by fluorescence microscopy: effect on mitosis and guidelines for non-invasive fluorescence microscopy. *The Plant Journal*, 36(2):280–290, 2003.
 51. RA Hoebbe, CH Van Oven, TWJ Gadella Jr, PB Dhonukshe, CJF Van Noorden, and EMM Manders. Controlled light-exposure microscopy reduces photobleaching and phototoxicity in fluorescence live-cell imaging. *Nature biotechnology*, 25(2):249, 2007.
 52. Stefan W Hell and Jan Wichmann. Breaking the diffraction resolution limit by stimulated emission: stimulated-emission-depletion fluorescence microscopy. *Optics letters*, 19(11):780–782, 1994.
 53. Mats GL Gustafsson. Nonlinear structured-illumination microscopy: wide-field fluorescence imaging with theoretically unlimited resolution. *Proceedings of the National Academy of Sciences*, 102(37):13081–13086, 2005.
 54. Eric Betzig, George H Patterson, Rachid Sougrat, O Wolf Lindwasser, Scott Olenych, Juan S Bonifacio, Michael W Davidson, Jennifer Lippincott-Schwartz, and Harald F Hess. Imaging intracellular fluorescent proteins at nanometer resolution. *Science*, 313(5793):1642–1645, 2006.
 55. Michael J Rust, Mark Bates, and Xiaowei Zhuang. Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (storm). *Nature methods*, 3(10):793, 2006.
 56. Nils Gustafsson, Siân Culley, George Ashdown, Dylan M Owen, Pedro Matos Pereira, and Ricardo Henriques. Fast live-cell conventional fluorophore nanoscopy with imagej through super-resolution radial fluctuations. *Nature communications*, 7:12471, 2016.
 57. Siân Culley, Kalina L Tosheva, Pedro Matos Pereira, and Ricardo Henriques. Srrf: Universal live-cell super-resolution microscopy. *The international journal of biochemistry & cell biology*, 2018.
 58. Elias Nehme, Lucien E Weiss, Tomer Michaeli, and Yoav Shechtman. Deep-storm: super-resolution single-molecule microscopy by deep learning. *Optica*, 5(4):458–464, 2018.
 59. Nicholas Boyd, Eric Jonas, Hazen P Babcock, and Benjamin Recht. Deeploco: Fast 3d localization microscopy using neural networks. *BioRxiv*, page 267096, 2018.
 60. AJ Nelson and ST Hess. Molecular imaging with neural training of identification algorithm (neural network localization identification). *Microscopy research and technique*, 81(9):966–972, 2018.
 61. Yu Li, Fan Xu, Fa Zhang, Pingyong Xu, Mingshuo Zhang, Ming Fan, Lihua Li, Xin Gao, and Renmin Han. Dibi: Deep learning guided bayesian inference for structure reconstruction of super-resolution fluorescence microscopy. *arXiv preprint arXiv:1805.07777*, 2018.
 62. Daniel Sage, Thanh-An Pham, Hazen Babcock, Tomas Lukes, Thomas Pengo, Jerry Chao, Ramraj Velmurugan, Alex Herbert, Anurag Agrawal, Silvia Colabrese, et al. Super-resolution fight club: A broad assessment of 2d & 3d single-molecule localization microscopy software. *bioRxiv*, page 362517, 2018.
 63. Stefan Bauer, Nicolas Carion, Peter Schüffler, Thomas Fuchs, Peter Wild, and Joachim M Buhmann. Multi-organ cancer classification and survival analysis. *arXiv preprint arXiv:1606.00897*, 2016.
 64. Jiawen Yao, Sheng Wang, Xinliang Zhu, and Junzhou Huang. Imaging biomarker discovery for lung cancer survival prediction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 649–657. Springer, 2016.

65. Oliver Dürr and Beate Sick. Single-cell phenotype classification using deep convolutional neural networks. *Journal of biomolecular screening*, 21(9):998–1003, 2016.
66. Mengjia Xu, Dimitrios P Papageorgiou, Sabia Z Abidi, Ming Dao, Hong Zhao, and George Em Karniadakis. A deep convolutional neural network for classification of red blood cells in sickle cell anemia. *PLoS computational biology*, 13(10):e1005746, 2017.
67. Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5967–5976. IEEE, 2017.
68. Aharon Azulay and Yair Weiss. Why do deep convolutional networks generalize so poorly to small image transformations? *arXiv preprint arXiv:1805.12177*, 2018.
69. Siân Culley, David Albrecht, Caron Jacobs, Pedro Matos Pereira, Christophe Leterrier, Jason Mercer, and Ricardo Henriques. Quantitative mapping and minimization of super-resolution optical imaging artifacts. *Nature methods*, 15(4):263, 2018.
70. Romain F. Laine, Kalina L. Tosheva, Nils Gustafsson, Robert D. M. Gray, Pedro Almada, David Albrecht, Gabriel T. Risa, Fredrik Hurtig, Ann-Christin Lindås, Buzz Baum, Jason Mercer, Christophe Leterrier, Pedro M. Pereira, Siân Culley, and Ricardo Henriques. Nanoj: a high-performance open-source super-resolution microscopy toolbox. *Journal of Physics D: Applied Physics*, 2019.
71. William Lotter, Gabriel Kreiman, and David Cox. Deep predictive coding networks for video prediction and unsupervised learning. *arXiv preprint arXiv:1605.08104*, 2016.
72. Daniel H Fisch, Artur Yakimovich, Barbara Clough, Joseph Wright, Monique Bunyan, Michael Howell, Jason Mercer, and Eva-Maria Frickel. An artificial intelligence workflow for defining host-pathogen interactions. *bioRxiv*, page 408450, 2018.
73. Joel Lehman, Jeff Clune, Dusan Misevic, Christoph Adami, Julie Beaulieu, Peter J Bentley, Samuel Bernard, Guillaume Belson, David M Bryson, Nick Cheney, et al. The surprising creativity of digital evolution: A collection of anecdotes from the evolutionary computation and artificial life research communities. *arXiv preprint arXiv:1803.03453*, 2018.