The advent of AI and deep learning in diagnostics and imaging

Machine learning systems have potential to improve diagnostics in healthcare and imaging systems in research

Philip Hunter

AI and machine learning (ML) in healthcare have a long history dating back to the 1980s that started with rule-based systems followed by hierarchical clustering and linear regression algorithms. However, neither the algorithms nor the computers themselves were yet sufficiently powerful to enable effective ML. Over the past 5 years though, advances in computational power combined with new algorithms based on neural network techniques have enabled enormous progress in ML, which is now impacting on many fields, including research and healthcare. Two fundamental capabilities of ML make it particularly interesting for diagnostics and research: the abilities to detect weak signals amidst noise and to enhance low-resolution images. However, there are also concerns about potential for abuse with potentially fatal consequences for applications in science and healthcare.

Blood-based cancer tests

On the diagnostics front, much research has focused on early detection of cancer based on blood tests, which has spawned several start-up firms that combine ML with molecular biology. Freenome, a Silicon Valley company, has developed a test based on multiple analytes, which, it claims, has higher levels of accuracy for detecting colorectal cancer than any existing test. “The underlying hypothesis at Freenome is that there is no single analyte solution to cancer screening”, said the firm’s CEO Gabe Otte. The theory is that combining multiple analytes, based on, say DNA, RNA and protein, can yield both greater sensitivity

and specificity than individual tests on their own.

Indeed, it is impossible to tune a single analyte test to maximize both sensitivity and specificity at the same time. Making the test as sensitive as possible comes at the expense of specificity, which increases the risk of false positives. This defeats one of the objects of a cancer test—to save money and make diagnosis simpler—because follow-up tests would still be required for confirmation. “Sensitivity and specificity are two separate parameters”, Otte clarified. “Most diagnostics are trying to optimize for both using just one degree of freedom, that is one analyte, and that is impossible because it’s a trade-off. But if you combine more than one analyte you can have both high specificity and sensitivity, because you can optimize for both with two degrees of freedom”.

“Two fundamental capabilities of ML make it particularly interesting for diagnostics and research: the abilities to detect weak signals amidst noise and to enhance low-resolution images.”

However, combining signals from multiple analytes to achieve multiple degrees of freedom is not trivial and that is where the machine learning comes in, Otte explained. “There’s many different ways of merging the signals from different analytes. The most traditional way is to do linear combination, taking the signals from each analyte and putting a different weight to it. For example,

you might decide you are less confident of the DNA signal than the RNA signal and so assign a lower weight to that. But these yield scores that when added up are ultimately meaningless”.

“... combining signals from multiple analytes to achieve multiple degrees of freedom is not trivial and that is where the machine learning comes in...”

The main challenge is establishing the right contribution of each signal and ensuring the measurements are orthogonal, that is, independent of each other. “With machine learning merging the signals can be done differently”, Otte said. “We can run computational trials, where the machine learning tries out the analytes, DNA, RNA and proteins, in different combinations and works out which yields the most accurate tests. This powerful iterative capability did not exist even 5 years ago”. Otte claimed that his team has achieved higher levels of sensitivity than most alternative liquid biopsies for colorectal cancer, even with a single analyte test based on DNA. “Most DNA applications today use Circulating Tumour (CT) DNA and these are extremely rare in early stage, so at stage one sensitivities of 30% to 40% are common. We’re using immune DNA recognizing immune cells attacking the tumour and achieved sensitivity of 83%”.

ML in diagnostics and treatment management

Liquid biopsies not only hold potential for early diagnostics of cancer but also for monitoring treatment so as to fine-tune the therapy. John Cassidy, cofounder and CEO of Cambridge Cancer Genomics (CCG), noted that every cancer is unique not just at the outset, but as it evolves during development and in response to treatment. CCG thus applies AI to analyse mutations in the circulating cell-free (cf) DNA from patients' blood samples both to monitor treatment and to screen for relapse. Cassidy commented though that AI or ML are not themselves magic bullets for treating cancer or any other diseases, but merely effective tools to combine disparate data to help in diagnosis and managing treatment. "So the power of AI/ML is very much in for example analysing a combination of scan data with genomics or liquid biopsy", Cassidy explained. "If applied in the correct way, then these techniques are an important tool in the future of medicine, but many more pieces of the puzzle are required".

"... an attack using GAN could add or remove evidence of aneurysms, heart disease, blood clots, infections, arthritis, cartilage problems, tumours in the brain, heart, or spine, and other cancers."

Cassidy gave the example of matching diagnostics with clinical trials databases to find studies that are relevant for a given patient. "At CCG.ai, we use neural networks to define mutations from a patient's tumour and then natural language processing to decipher information on clinical trials.gov", he said. "From there, we can match patients, and their genomic signature, to appropriate clinical trials. This would be an example of scalable precision or personalised medicine built on machine learning".

As early diagnosis of solid tumours can often be achieved via X-ray, MRI or ultrasound scanning, AI and especially ML are already making an impact by helping radiologists to find patterns and structures in the image data that would indicate a tumour. Whether such patterns reside in molecular data or scanned images is almost

irrelevant as far as the underlying technology is concerned.

Potential for abuse

However, these applications come with a significant risk of fraud and abuse. In 2019, an Israeli research group published its results after an attack on a hospital's picture archiving and communication system with the aim of installing malware and tampering with 3D medical images autonomously [preprint: 1]. The team used an approach to ML called GAN (generative adversarial network) whereby two separate neural network-based ML systems challenge each other. Under GAN, one ML system creates fake images while the other attempts to distinguish these from genuine ones. This process in which the two ML systems learn from each other and become more finely tuned to their task enables the second system to better recognize genuine images and the first to become very good at faking.

The Israeli team demonstrated that an attack using GAN could add or remove evidence of aneurysms, heart disease, blood clots, infections, arthritis, cartilage problems, tumours in the brain, heart, or spine, and other cancers. The authors suggested various motives for such attacks: such as influencing the outcome of an election or toppling a political figure by prompting a serious health diagnosis. It also noted that, although no such attack had taken place, clinics and hospitals had been hit by numerous data breaches and interruptions in medical services as a result of hacks during 2018. "I believe that this threat will become more apparent as time goes on, over the next few years", said Yisroel Mirsky, project manager at Ben-Gurion University of the Negev, a public research university in Israel and lead author on the paper. "That is especially the case with AI being so accessible and given the power of state actors. Currently, the insurance fraud scenario is more likely to be an immediate concern".

However, countermeasures are available, Mirsky added. "The best way is to enable digital signatures, having all scanners sign each scan and have the signature verified by the viewing applications. Then if tampering occurs between the end-points the attack will be discovered immediately". He also suggested that hospitals should deploy public key infrastructure technology within their networks to secure digital signatures.

Breaking the Abbe limit

As with medical imaging, AI techniques are being increasingly applied in optical microscopy, mostly to overcome the resolution limit of light-microscopy systems without the need for sophisticated and expensive high-resolution microscopy. The major impetus is the fact that, although optical technologies can never match the resolution of electron microscopy, it enables the observation of living biological processes in real time. This in turn raises other challenges, such as avoiding damage of the sample caused by the necessary illumination and the reliance of fluorescence markers to label specific proteins or structures.

The resolution of conventional optical microscopy is limited by diffraction as light rays are bent around the edges of the lens system. This determines the so-called Abbe limit: the shortest possible distance between two objects that can be distinguished by a viewer. As electron microscopy uses X-rays with a shorter wavelength to illuminate the sample, it can therefore differentiate between much smaller objects and details.

"... ML creates the possibility of operating purely on image data from lower-cost equipment, even from mobile-phone cameras."

However, the Abbe limit was recast by Richard Feynman on the basis of time: the new rule being that two different point sources can be resolved only if the time taken for rays from each of them to reach the eyepiece differs by more than one period of the wavelength used (http://www.feynmanlectures.caltech.edu/I_30.html). This opened the door to a number of methods under the banner of super-resolution microscopy that exploited either the new rule in various ways or other aspects of physics [2]. But ML creates the possibility of operating purely on image data from lower-cost equipment, even from mobile-phone cameras.

All ML techniques to improve the resolution of optical microscopy images rely on the same underlying principle that low-resolution images contain sufficient information to generate higher resolution images. The point is that ML can be trained on pairs of

images where one is high and the other is low resolution to make a prediction of what the high-resolution image should look like. This prediction is then compared with the actual high-resolution image and the difference calculated as the “error back propagation”. The model then adjusts and has another go. This iterative process is continued until the prediction is as close to the real high-resolution image as possible.

There are various approaches being pursued to extract high-resolution data from low-resolution images. “Our method is based on optical acquisition, rather than relying on other algorithms to synthesize it”, explained Yair Rivenson, a post-doctoral fellow in the Ozcan Research group at the University of California, Los Angeles, California (UCLA). Unlike most other statistical methods for achieving super-resolution, their approach is driven by real data from actual high-resolution microscopic imaging to train the ML system based on a deep neural network. “This gives a few advantages in terms of enabling the network to learn data transformations that do not depend on simplifying assumptions on the imaging system and sample characteristics as some of the other techniques do”, Rivenson said.

One such advantage lies in the ability to perform cross-modality imaging from different types of optical microscopes to yield detail at higher resolution than either on its own. “We have demonstrated transformation between two different kinds of microscopes, for example, confocal to stimulated-emission-depletion microscopy and total internal reflection microscopy to structured illumination total internal reflection microscopy, in both cases beating the optical limit of diffraction”, Rivenson added [3].

Apart from overcoming the diffraction limit, ML is also applicable to a wider range of optical microscopes. “The potential is democratizing of microscopic imaging, by enabling it to do measurements with lower-cost equipment, while still retaining the performance of the high-resolution microscopy system”, Rivenson explained. He referred to recent work showing that the technique can even work with portable field microscopes combined with standard smartphone cameras [4]. Rivenson and colleagues succeeded in applying ML to correct distortions and enable high-resolution, colour-corrected images that, he claimed, match the performance of conventional benchtop microscopes with high-end objective lenses.

Augmented reality in microscopy

Machine learning is also being applied for analysing the components of microscopy images in various ways, to identify regions of specific interest or to spot features human observers might have missed. This is further enhanced by augmented reality (AR) techniques that combine digitally created visual components with real images. Google is among the leaders on this front with its Augmented Reality Microscope (ARM) platform that applies ML and AR to standard brightfield microscopy. The immediate goal is to make the system work in real time to analyse images quickly for applications in diagnostics and research. “The primary enhancements are computational optimizations to speed up the application of the machine learning models, or inference, and enable real-time usage”, explained Cameron Po-Hsuan Chen, software engineer and researcher at Google Brain, Google’s deep learning and AI team. “In our study, these enhancements improved responsiveness from one inference every 2-s to 5 inferences per second. Additional improvements shown by new data under review speeds this up further to 30 inferences per second”.

Google’s ambition was to calculate the ML output in real time so that it can be integrated with the image as it appears in the eyepiece. As well as superimposed graphics, this can include pointers to specific areas or the results of calculations such as cell counts performed on the image’s components. “Other potential applications in life science include AI-based stain quantification, mitosis counting and cell counting”, Chen added. “Beyond the life sciences, the ARM can potentially be applied to other microscopy applications such as material characterization in metallurgy and defect detection in electronics manufacturing” [preprint: 5].

Another application of ML in optical microscopy is the ability to image and draw data from large areas, which has been a major constraint for super-resolution techniques. A team at the Imaging and Modelling Unit of the Département Biologie Cellulaire et Infections at the Institut Pasteur in France has developed an algorithm called ANNA-PALM to improve temporal resolution, and reduce the time needed to acquire a super-resolution image [6]. “ANNA-PALM is based on a previously proposed image-to-image transformation approach called pix2pix, which itself is based on so-called U-nets, a special type of neural

network, and conditional generative adversarial networks, another recent deep learning technique that uses two neural networks trained one against the other”, said Christophe Zimmer, head of that imaging unit. “As we showed in our paper, ANNA-PALM can be used to image thousands of cells at high resolution rather than just a handful of cells. This should make it easier to detect rare phenotypes, for example subpopulations of cells in which microtubule architecture is somehow disrupted. In general, we think that ANNA-PALM will be useful in high-throughput image-based screening of drugs or gene knock-out screens”.

.....
“Another application of ML in optical microscopy is the ability to image and draw data from large areas, which has been a major constraint for super-resolution techniques.”

“Electron microscopy can indeed achieve much better resolution than optical microscopy, but usually does not allow us to highlight specific molecules”, Zimmer explained the importance of these recent developments in optical microscopy. “By contrast, with fluorescence microscopy we can visualize specific molecules of interest, say tubulin or actin or a particular nuclear pore component”.

Environmental monitoring

This flexibility of optical microscopy can still be taken a stage further by dispensing with the need for a lens system at all, which opens the door to a variety of other applications including environmental monitoring. IBM develops such devices it calls AI microscopes that comprise digital sensors attached to a computing board.

“We put the sample directly on the digital sensor, and then shine the point light on the sample”, explained Simone Bianco, Research Staff Member in IBM’s department of Industrial and Applied Genomics. “The camera will then capture the shadow for example of any creature swimming in the sample. The advantage of our device is that, since there are no optics, there is no focus, and we can clearly see every creature regardless of their swimming pattern. This would not be possible with a classical microscope, because

every time a creature falls off the focal plane, it becomes blurry. Once videos and images are acquired, we use computer vision and digital holography to increase the resolution and extract the important features which are required to train our AI”.

IBM employed several techniques, including both supervised learning based on labelled data sets and unsupervised learning operating on raw data. “The difference between supervised and unsupervised learning is that supervised learning algorithms require large and accurate annotation to be trained, while unsupervised algorithms do not require annotation, but require additional steps in order to be as accurate and assign labels to what is in the samples”, Bianco said. “In many scientific fields, annotation is hard and most times very costly to obtain, and so we are pushing for an unsupervised AI to power our microscopes”.

Bianco anticipates their use for continuous monitoring of aquatic environments. “The aim is to use our AI to learn plankton shape and behaviour and use it as a sensitive and accurate early biosensor of contamination of our waters”, he said.

Such applications are just the first ones by which AI and ML are being employed in clinical practice and life science research. There will surely be more to come but so there will be more concerns about the potential of ML to generate fake data or images.

References

1. Mirsky Y, Mahler T, Shelef I, Elovici Y (2019) CT-GAN: malicious tampering of 3D medical imagery using deep learning. *arXiv* <https://arxiv.org/abs/1901.03597> [PREPRINT]
2. Cremer C, Masters BR (2013) Resolution enhancement techniques in microscopy. *Eur Phys J H* 38: 281–344
3. Wang H, Rivenson Y, Jin Y, Wei Z, Gao R, Günaydin H, Bentolila LA, Kural C, Oxcan A (2018) Deep learning enables cross-modality super-resolution in fluorescence microscopy. *Nat Methods* 16: 103–110
4. Rivenson Y, Koydemir HC, Wang H, Wei Z, Ren Z, Günaydin H, Zhang Y, Göröcs Z, Liang K, Tseng D et al (2018) Deep learning enhanced mobile-phone microscopy. *ACS Photonics* 5: 2354–2364
5. Chen PH, Gadepalli K, MacDonald R, Liu Y, Nagpal K, Kohlberger T, Dean J, Corrado GS, Hipp JD, Stumpe MC (2018) Microscope 2.0: an augmented reality microscope with real-time artificial intelligence integration. *arXiv* <https://arxiv.org/abs/1812.00825> [PREPRINT]
6. Ouyang W, Aristov A, Lelek M, Hao X, Zimmer C (2018) Deep learning massively accelerates super-resolution localization microscopy. *Nat Biotechnol* 35: 460–468