

Editorials

Harnessing the power of intelligent machines to enhance primary care

Mr Jones sits down in front of you and removes from his bag a sheaf of printed web pages. Trying to appear composed, you studiously ignore the stack of papers and look straight at him. *'How can I help you today?'*

'Well, doctor, it's just that I was looking on the internet and they suggested I come in and see you ...' [he looks bashful] *'... it's just I've been having difficulty swallowing and the weight's been dropping off me. The internet suggested it might be ...'* [he hesitates] *'... something nasty.'*

Two weeks later you open your correspondence. This time, the internet was right.

Searching the internet for health advice is undoubtedly fraught with potential complications as well as rewards, but have you ever stopped to think about the technology behind these searches and what other implications it might have for your practice?

MACHINE LEARNING

Only 9% of the UK population have heard of machine learning, but 89% recognise at least one of its uses.¹ From fraud detection on your credit card to voice recognition software such as Siri, we are surrounded by intelligent machines. Each time a search term is entered into Google it uses complex algorithms to determine the most appropriate results. These rules are not programmed explicitly by humans but developed iteratively over time by the machine itself. The machine is learning from our behaviour — which results we choose to click on, how long we stay on an individual page — and is constantly refining its own search algorithms as a result of these feedback loops.

Unsurprisingly, it is difficult to mimic the action of the human brain's 86 billion neurones to any significant degree, and the much hyped artificial intelligence that might produce robotic doctors is a long way off. The closest technology has come so far is in the use of artificial neural networks, composed of a complex system of inputs and outputs (see Figure 1). The repeated activation of connections strengthens their bond, applying differing weights to the various contacts. As hidden layers are added, the networks' complexity intensifies, their ability




Figure 1. Diagrammatic representation of artificial neural networks.

to learn increases, and they become known as 'deep' neural networks. The deeper these networks, the more difficult it becomes for those on the outside to understand what is happening inside the 'black box' of hidden layers. Although this is a legitimate concern, it is important to remember that there is precedent in medicine for improving health without fully understanding the mechanism of action of the tools being used, for instance, lithium in bipolar disorder.²

In evaluating machine learning we must remember that the comparator is real-world clinical care, performed by fallible humans supported by general guidelines that can be difficult to apply to an individual patient. Of course, over-reliance on data can lead to errors, as it does not take into account the nuances of the consultation.³ Concerns about deskilling of professionals through use of machine learning can be weighed against the potential to mechanise mundane repetitive work, reduce errors, and free up time for patient interaction — the true heart of general practice.

Rather than concentrating on the negative aspects of this technology, we should focus on how machine learning could enhance primary care; specifically here we explore the impact it might have on prevention, diagnosis, and management.

PREVENTION

Three examples, from very different spheres of medicine, illustrate the potential impact of machine learning in the prevention of disease and the promotion of health. Reinforcement learning (active learning from observed behaviours) can develop personalised text messages that improve levels of physical activity, recognising that individuals respond to encouragement in different ways.⁴

Personalised parameters are also generated by Sention™, a system using machine learning to prevent hospital admissions. In patients at risk of avoidable admission it collects information from simple wearable sensors to detect changes suggesting deterioration and then contacts a clinician. This type of technology could revolutionise the community-based care of individuals with serious and complex health needs.

Predicting disease can also aid personalised prevention; recently a neural network improved accuracy of prediction of cardiovascular events over 10 years by 3.6% compared with standard criteria.⁵ The top 10 predictor variables for the neural network were very different from those of the standard algorithm and included severe mental illness and corticosteroid use,

“Machine learning is not only being used by doctors, but also by patients.”

factors that have since been introduced into the latest QRISK@3 calculator.⁶

DIAGNOSIS

A ground-breaking study showed that machine learning can classify skin cancers with the accuracy of a dermatologist.⁷ The neural network was trained using 129 450 images of 2032 different skin diseases; the only inputs were the pixels from the images and the disease labels. The resultant system was able to classify melanomas versus benign naevi at least as accurately as dermatologists. There is great scope for similar technology in other areas of health care where rich data are central to diagnosis, such as detecting diabetic eye disease from retinal photographs.⁸

It is not only in classification of disease that machine learning has been applied; it has also been utilised in case identification. A systematic review showed that analysis of free text in healthcare records by machine learning techniques increases case detection significantly compared with analyses based on recorded codes.⁹ Embedding software for these algorithms into our clinical systems, providing automatic surveillance partnered with continual learning, could help identify cases across many areas of disease. However, given the relatively low incidence of new disease in primary care,¹⁰ it is likely that prompts will have a high false-positive rate, so software developers and clinicians will have to work together to avoid alert fatigue.¹¹

MANAGEMENT

Mobile apps are beginning to utilise the vast amount of ‘exhaust data’ available from your mobile phone. Ginger.io uses information such as the number of text messages sent or how often you leave your house to monitor symptoms of depression. Using machine learning, the algorithm identifies usual personal routines and alerts the user (or pre-arranged healthcare providers) when there is deviation. This approach of analysing divergence from normal behaviour can also be applied to medication management systems. One such system, which learns to detect statistical outliers, can constantly survey electronic records for potential drug errors,

for instance, when a patient on NSAIDs develops a high creatinine, and provides alerts for the clinician.¹²

Machine learning is not only being used by doctors, but also by patients. Ada is one of many apps exploiting it to provide relevant information to patients on the likely cause of their symptoms and subsequent management accompanied by links to NHS Choices. This technology does not aim to replace primary care management of chronic disease, nor provide emergency care, but it does have the potential to enhance self-care.

THE FUTURE

Machine learning already permeates many areas of our life. Health care is not immune from this and nor should we hope to be, as this technology has enormous potential to improve patient outcomes. Instead of being fearful of machines usurping our roles, we should embrace their ability to enhance the care we provide and appreciate the time they can create for doctors to focus on what makes them truly indispensable — human contact.

Sarah F Moore,

Academic Clinical Fellow and GP trainee, University of Exeter Medical School, Exeter.

William Hamilton,

Professor of Primary Care Diagnostics, University of Exeter Medical School, Exeter.

David J Llewellyn,

Senior Research Fellow, University of Exeter Medical School, Exeter.

Provenance

Commissioned; externally peer reviewed.

Funding

Sarah F Moore is a locally funded Academic Clinical Fellow in a position approved by the UK National Institute for Health Research (NIHR). David J Llewellyn is supported by the Halpin Trust and he and William Hamilton are supported by the UK NIHR Collaboration for Leadership in Applied Health Research and Care South West Peninsula (PenCLAHRC). The views expressed in this article are those of the authors and not necessarily those of the NHS, the NIHR, or the Department of Health.

DOI: <https://doi.org/10.3399/bjgp17X693965>

ADDRESS FOR CORRESPONDENCE

Sarah F Moore

University of Exeter Medical School, College House, St Luke's Campus, Magdalen Road, Exeter, EX1 2LU, UK

E-mail: smoore15@nhs.net

REFERENCES

1. Royal Society. *Machine learning: the power and promise of computers that learn by example*. London: Royal Society, 2017.
2. Malhi GS, Tanious M, Das P, *et al*. Potential mechanisms of action of lithium in bipolar disorder: current understanding. *CNS Drugs* 2013; **27**(2): 135–153.
3. Cabitza F, Rasoini R, Gensini GF. Unintended consequences of machine learning in Medicine. *JAMA* 2017; **318**(6): 517.
4. Hochberg I, Feraru G, Kozdoba M, *et al*. Encouraging physical activity in patients with diabetes through automatic personalized feedback via reinforcement learning improves glycemic control. *Diabetes Care* 2016; **39**(4): e59–e60.
5. Weng SF, Reips J, Kai J, *et al*. Can machine-learning improve cardiovascular risk prediction using routine clinical data? *PLOS One* 2017; **12**(4): e0174944.
6. Hippisley-Cox J, Coupland C, Brindle P. Development and validation of QRISK3 risk prediction algorithms to estimate future risk of cardiovascular disease: prospective cohort study. *BMJ* 2017; **357**: j2099.
7. Esteva A, Kuprel B, Novoa RA, *et al*. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 2017; **542**(7639): 115–118.
8. Gulshan V, Peng L, Coram M, *et al*. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA* 2016; **316**(22): 2402.
9. Ford E, Carroll JA, Smith HE, *et al*. Extracting information from the text of electronic medical records to improve case detection: a systematic review. *J Am Med Inform Assoc* 2016; **23**(5): 1007–1015.
10. Buntinx F, Mant D, Van den Bruel A, *et al*. Dealing with low-incidence serious diseases in general practice. *Br J Gen Pract* 2011; DOI: <https://doi.org/10.3399/bjgp11X548974>.
11. Nanji KC, Slight SP, Seger DL, *et al*. Overrides of medication-related clinical decision support alerts in outpatients. *J Am Med Inform Assoc* 2014; **21**(3): 487–491.
12. Schiff GD, Volk LA, Volodarskaya M, *et al*. Screening for medication errors using an outlier detection system. *J Am Med Inform Assoc* 2017; **24**(2): 281–287.