

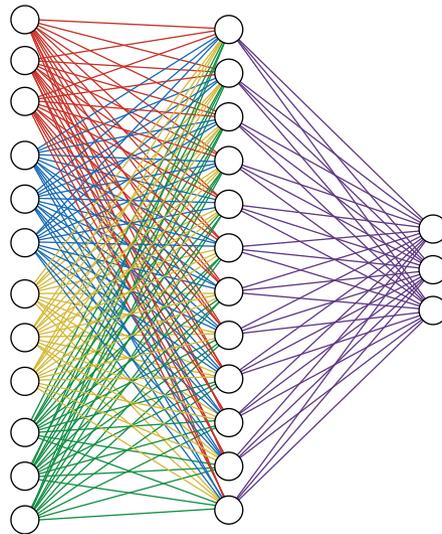
# Auspicious machine learning

The accelerating diagnostic power of deep learning will soon empower physicians.

Perception is not uniquely a human ability. Computers running deep convolutional neural networks and other types of machine-learning algorithm can also recognize objects in pictures and discriminate subtle differences in otherwise similar patterns. Today's phones make use of machine learning to, for example, make sense of the owner's words and accurately translate text between languages — tasks that only a few years ago were exclusively done by humans. Beyond allowing mobile-phone users to offload easy tasks (such as tagging faces in photos) and perform jobs that require vast amounts of knowledge (such as finding the least busy route to a destination), machine learning will soon help us to live healthier and longer.

Deep convolutional neural networks loosely imitate the multilayered organization of the visual cortex. They can be thought of as a stack of interconnected layers of nodes, each node carrying out a computation: a predetermined smooth mathematical function that takes in the dot product (an algebraic operation) of the inputs and the weights of the connections carrying the inputs. On the whole, neural networks transform input volumes of data (for example the pixels in an image) on the basis of the weights of each connection, which are determined a priori by training the network with known input and output volumes. Essentially, training involves finding the weights that minimize the difference between the known inputs and the desired outputs.

The deeper the convolutional network — that is, the larger the number of layers — the greater the levels of abstraction in the representation of the data, and the more accurate its 'perception'. Modern convolutional networks consist of 10–20 layers and over 100 million nodes, weights and biases (the values of nodes when the inputs are zero), and can be used with techniques that prevent overfitting — that is, the fitting of any noise present in the training dataset. Naturally, neural networks are hungry for training data — the bigger the dataset, the better they will perform when presented with new data. For example, about 130,000 clinical images of skin lesions<sup>1</sup> and of retinal images<sup>2</sup> were recently used to train convolutional networks to classify skin cancers<sup>1</sup> and



detect diabetic retinopathy<sup>2</sup>, with accuracy comparable to expert dermatologists and ophthalmologists. Yet, can convolutional networks help with rare diseases, for which data are scarce?

For congenital cataracts — a rare disease that leads to clouded vision in children — the answer is yes. As reported by Long *et al.* (article no. 0024) in this issue, a neural network with five convolutional layers trained with 886 ocular images was able to perform just as well as individual expert ophthalmologists in diagnosing the disease, in determining its severity (according to three parameters for eye-lens opacity) and in suggesting treatment (surgery or follow-up). Notably, the study was carried out in real-world settings: a clinical trial run by a network of hospitals as well as a 'needle in a haystack' test with a realistic ratio of the incidence of congenital cataracts (~1% of the general population). Moreover, the machine-learning algorithm was implemented in a bespoke cloud-based platform for multihospital collaboration. In a News & Views (article no. 0032), Wang and Shen note that the software platform could improve the management of rare diseases, especially in care centres with insufficient resources.

Machine learning can also help to diagnose brain tumours. Orringer *et al.* (article no. 0027) report in this issue that a multilayer perceptron (another type of deep machine learning) is able to categorize brain-tumour subtypes from images of brain

tissue, rapidly generated from stimulated Raman scattering microscopy of fresh surgical samples, that visually resemble those images acquired from the stained, frozen or fixed tissue sections processed by pathology laboratories. By using such 'virtual histology', neuropathologists diagnosed brain tumours as accurately (over 90%) as when using conventional tissue sections. Stimulated Raman microscopy may thus rapidly provide key diagnostic information to surgeons after tumour resection, as noted by Figueiredo and Rutka in a News & Views (article no. 0033).

The utility of machine learning goes beyond medical diagnostics. As demonstrated by Farina *et al.* (article no. 0025), a support vector machine (a machine-learning algorithm for sorting data) is able to categorize into discrete classes the activity of spinal motor neurons surgically redirected into muscles of the chest (known as targeted muscle reinnervation) in upper-limb amputees. The classified data can be mapped into control commands for the eventual control of a prosthetic. In a News & Views (article no. 0034), Ortiz-Catalan remarks that the decoding of motor-neuron activity should also enable the monitoring of the progression of muscle reinnervation.

Ultimately, computer-aided classification or diagnosis should reduce inevitable fallibilities and inherent diagnostic variabilities in healthcare (especially in routine tissue pathology), provide guidance (above all, in settings where expert physicians are not available), and speed up decision-making, in particular in the operating room. The seemingly vast and promising opportunities that machine learning will bring to healthcare will also advance precision medicine and disease management at the level of individual patients.

Soon enough, today's machine-learning algorithms will resemble what the mobile phones of the 1990s are to today's smartphones. Software powered by machine learning will be able to drive cars, make the best personal concierges and discover new drugs. A physician's touch will however remain unique. □

## References

1. Esteva, A. *et al.* *Nature* <http://dx.doi.org/10.1038/nature21056> (2017).
2. Gulshan, V. *et al.* *JAMA* **316**, 2402–2410 (2016).