



50 Years Ago

In spite of the enlightened attitudes of many countries, the opinion still prevails in Britain that engineering is not a suitable career for women. In France, one engineer in twenty-eight is a woman, and in Syria one in fourteen, while in Russia the figure is one in three. But in Britain, only one engineer in five hundred is a woman ... Last week, the Women's Engineering Society ... celebrated its fiftieth anniversary. At the same time "Women in Engineering Year" was launched in a concerted effort to break down the prejudice ... against women taking their place in a profession which needs as many eager recruits as it can get. Conferences, exhibitions, lectures and visits have been organized throughout Britain to demonstrate that engineering does not consist entirely of heavy and dirty work requiring massive physical stamina, and that women have a valuable part to play.

From *Nature* 15 March 1969

100 Years Ago

Of late years much attention has been given to the remarkable power of charcoal to absorb gases of all kinds, and during the war extensive use has been made of this property in the construction of masks for removing noxious gases from the air inhaled by the wearer ... I should like to remind readers of *Nature* that the first practical application of charcoal for such purposes was made by Dr. John Stenhouse, lecturer in chemistry at St. Bartholomew's Hospital. In 1854 Stenhouse devised a charcoal respirator consisting of a perforated zinc case filled with granular wood charcoal, and adapted to fit over the mouth and nose. Respirators of this kind were in use by nurses and dressers in St. Bartholomew's ... down to the time when Lister's antiseptic system rendered such protection from the offensive emanations of sores unnecessary.

From *Nature* 13 March 1919

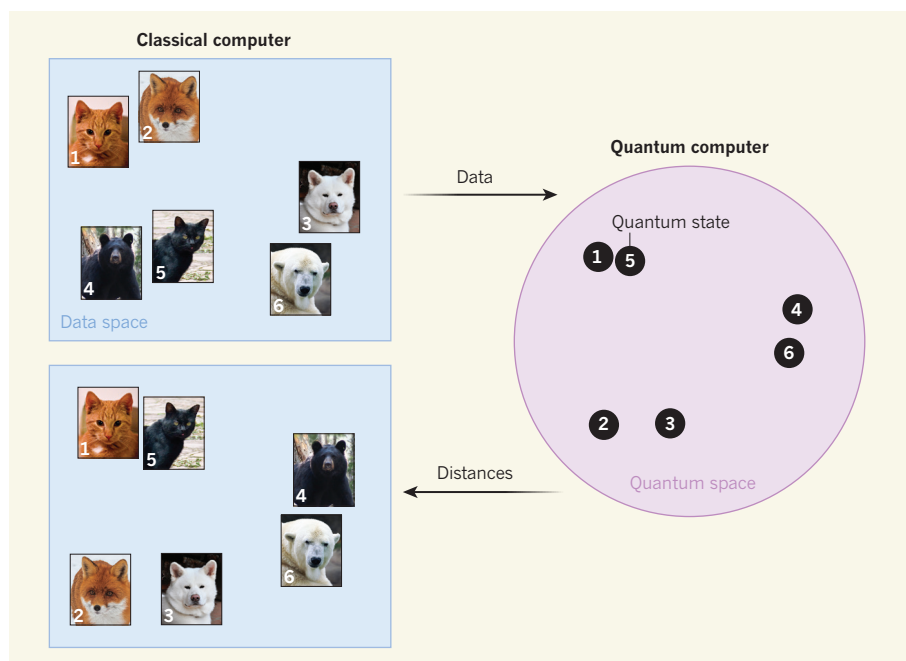


Figure 1 | Quantum-enhanced machine learning. Havlíček *et al.*¹ demonstrate how quantum computers could improve the performance of machine-learning algorithms. In this simple illustration, a conventional (classical) computer uses machine learning to classify images of animals. Images whose pixels contain similar colours are positioned close together in data space. The classical computer sends these data to a quantum computer that maps each of the images to a particular quantum state in a space of such states. Images that are close together in data space, but are different in content, are represented by states that are far apart in quantum space. The quantum computer sends the distances between the quantum states to the classical computer to improve the image classification.

machine learning by defining which data points are similar to each other and which are not. Mathematically speaking, similarity is a distance in data space — that is, a distance between the representations of data points as numbers. Similar images are assumed to have similar content, and distances between data points can be crucial in machine learning. But defining similarities is not as straightforward as it sounds. For example, what is the distance in data space between two images if derived on the basis of the amount of red in each image?

Kernel theory showed that many definitions of similarity in data space are mathematically equivalent to a simple measure of similarity in a much larger, possibly infinitely large, space (Fig. 1). Consequently, every time two images are compared, the images are implicitly mapped to a representation in a huge space, and a simple similarity is computed. No ordinary computer can calculate this large representation explicitly. But perhaps a quantum computer can? Because quantum computers carry out computations in extremely large spaces, what happens if data are mapped into the space that is inhabited by quantum states?

Havlíček *et al.* and my research team³ recognized this potentially powerful link between machine learning and quantum computing at roughly the same time. Remarkably, both groups proposed essentially the same two strategies for using the idea to design quantum algorithms for machine learning. The first strategy makes only minimal use of

the quantum computer, as a mere hardware addition to a conventional machine-learning system: the quantum device returns similarities when given two data points. The second strategy carries out the actual learning on the quantum computer, aided by the classical one.

A key contribution from Havlíček *et al.* is that they implemented the two strategies in a proof-of-principle experiment on a real quantum computer: one of IBM's quantum chips. Despite the inflated claims of some news reports, anyone who has tried quantum computing in the cloud knows that collecting meaningful data from these devices is notoriously difficult, owing to the high levels of experimental noise in the computation. This is probably why the authors' experiment is stripped to its bare bones, in some people's view, maybe too much. The quantum space has only four dimensions, because the set-up uses two quantum bits (qubits) of IBM's smallest, five-qubit chip — at a time when the IBM cloud service already offers access to a 20-qubit device. The data set is likewise hand-engineered in such a way that it is simple to analyse in this four-dimensional space.

Nevertheless, Havlíček and colleagues' work presents an intriguing proof-of-principle demonstration of a potentially revolutionary way of using quantum computers for machine learning. After many studies offering various attempts to mould the much more popular artificial neural networks into quantum computing, kernel methods provide a refreshingly