where bone tissue forms) at the ends of bones soon after birth. Newton *et al.* investigated the mammalian target of rapamycin complex 1 (mTORC1) pathway, which has been reported to regulate stem-cell function[3]. They found that chondrocyte-specific activation of mTORC1 signalling leads to a shift from asymmetric to symmetric stem-cell divisions, and consequently to an increased number of stem cells in the resting zone. These observations strongly support a role for mTORC1 in regulating the self-renewal potential of resting-zone stem cells.

Both groups also analysed the role of the protein Indian hedgehog (Ihh), a member of the Hedgehog family of growth factors that is expressed in early-differentiated hypertrophic chondrocytes. Ihh has been shown to induce the expression of PTHrP in resting-zone chondrocytes, which in turn inhibits the premature initiation of hypertrophy in proliferating cells[4]. Additionally, both Ihh and PTHrP activate chondrocyte proliferation[3].

Newton *et al.* and Mizuhashi *et al.* provide evidence that the inhibition of Hedgehog signalling reduces the length of chondrocyte columns. Newton and colleagues also observed increased proliferation and the expression of genes targeted by Hedgehog proteins in resting-zone cells after activation of the Hedgehog pathway. These findings suggest that Hedgehog signalling has a role in controlling the stem-cell character of resting-zone cells.

However, given that Ihh regulates PTHrP expression directly, the observed changes in chondrocyte-column length and cell proliferation might also be a consequence of altered PTHrP levels. Furthermore, when Newton *et al.* inhibited Hedgehog signalling and activated the mTORC1 pathway simultaneously, some stem cells moved from the resting zone into the proliferating zone without differentiating into flat cells. Together, these observations support a role for Ihh in regulating stem-cell proliferation rather than stem-cell identity. Given that the interaction between Ihh and PTHrP signalling is complex, it will be challenging to distinguish clearly between the roles of Ihh as a regulator of stem-cell proliferation, PTHrP expression and the induction and maintenance of 'stemness'.

The model of how cartilage is replaced by bone has changed substantially in the past few years. Previously, hypertrophic chondrocytes were thought to die and then be replaced by bone-forming cells called osteoblasts. However, more recent fate-mapping studies have shown that a fraction of hypertrophic chondrocytes differentiate into bone-forming osteoblasts or long-lived stem cells and progenitor cells of the bone-marrow stroma[5-7]. Mizuhashi *et al.* now demonstrate that cells that are descendants of resting-zone stem cells contribute to the bone-marrow stroma. Therefore, such stem cells seem to follow an unusual path of differentiation, transforming from stem cells of the chondrocyte lineage into differentiated chondrocytes, and then into multilineage stem cells of the bone-marrow stroma.

Future investigations should clarify how many of the postnatal bone-marrow stem cells descend from resting-zone stem cells, and whether these postnatal cells differ functionally from other bone-marrow cells. Given that bone-marrow-derived skeletal stem cells are required for bone turnover and fracture repair throughout a person's life, deciphering the specific features of the chondrocyte-derived population will be of high clinical relevance.

The identification of a growth-plate-specific skeletal stem cell is an important step towards understanding human skeletal growth and associated diseases, but many questions remain. Follow-up studies need to determine which mechanisms besides Hedgehog and mTORC1 signalling induce and maintain the stem-cell character of these cells, which type of embryonic chondrocyte evolves into a resting-zone stem cell, and how the induction of that process is linked to the formation of secondary ossification centres.

Further studies also need to clarify how the differentiation of stem cells in the resting zone is regulated, and which components of the chondrocyte-specific extracellular matrix (the network of proteins and sugar molecules that surrounds cells) are required to generate a stem-cell niche. Finally, given that some hypertrophic chondrocytes differentiate into osteoblasts and bone-marrow stem cells, whereas others die[5-7], it is tempting to ask whether the fate of hypertrophic cells is already determined by the distinct subtypes of resting-zone stem cell from which they originate. ∎

Manuela Wuelling *and* **Andrea Vortkamp** *are at the Centre for Medical Biotechnology, Department of Developmental Biology, University Duisburg-Essen, 45141 Essen, Germany.*
e-mails: manuela.wuelling@uni-due.de; andrea.vortkamp@uni-due.de

1. Newton, P. T. *et al. Nature* **567**, 234–238 (2019).
2. Mizuhashi, K. *et al. Nature* **563**, 254–258 (2018).
3. Ito, K. & Suda, T. *Nature Rev. Mol. Cell Biol.* **15**, 243–256 (2014).
4. Ohba, S. *J. Dev. Biol.* **4**, E20 (2016).
5. Park, J. *et al. Biol. Open* **4**, 608–621 (2015).
6. Yang, L., Tsang, K. Y., Tang, H. C., Chan, D. & Cheah, K. S. *Proc. Natl Acad. Sci. USA* **111**, 12097–12102 (2014).
7. Zhou, X. *et al. PLoS Genet.* **10**, e1004820 (2014).

This article was published online on 27 February 2019.

INFORMATION SCIENCE

Machine learning in quantum spaces

Ordinary computers can perform machine learning by comparing mathematical representations of data. An experiment demonstrates how quantum computing could use quantum–mechanical representations instead. SEE LETTER P.209

MARIA SCHULD

Machine learning and quantum computing have their staggering levels of technology hype in common. But certain aspects of their mathematical foundations are also strikingly similar. On page 209, Havlíček *et al.*[1] exploit this link to show how today's quantum computers can, in principle, be used to learn from data — by mapping data into the space in which only quantum states exist.

One of the first things one learns about quantum computers is that these machines are extremely difficult to simulate on a classical computer such as a desktop PC. In other words, classical computers cannot be used to obtain the results of a quantum computation. The reason is that a lot of numbers are required to describe each internal step of the computation. Consider the multi-step procedure that many people learn at school for dividing large numbers. If this were a quantum computation being simulated on a classical computer,

every intermediate step could easily need more numbers to describe it than there are atoms in the observable Universe.

The state of a quantum system when described by a collection of numbers is known as a quantum state. And if a quantum state is associated with many values, it is said to 'live' in a large space. For certain quantum computers that are based on continuous variables, such spaces can even be infinitely large.

Machine learning, by comparison, analyses data that live in much smaller spaces — that is, the data are described by many fewer values. For example, a photograph that contains one million pixels records just three million numbers to describe the amount of red, green and blue in each pixel. A prominent task of machine learning could be to guess the content of the image, or to produce similar images. However, a well-established theory in machine learning called kernel methods[2] treats data in a way that has a similar feel to how quantum theory deals with data.

In a nutshell, kernel methods carry out

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
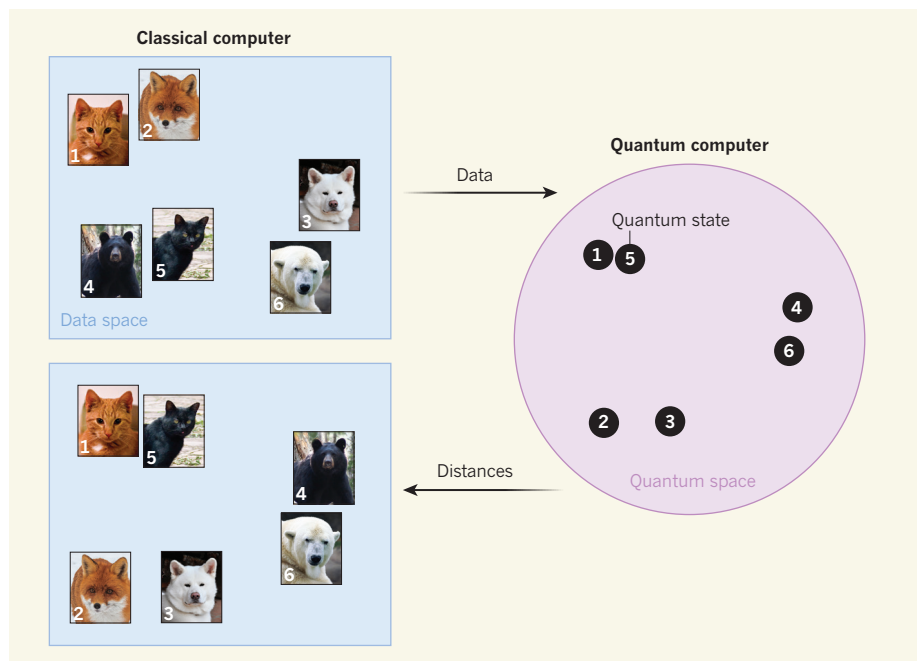
Figure 1 | Quantum-enhanced machine learning. Havlíček *et al.*[1] demonstrate how quantum computers could improve the performance of machine-learning algorithms. In this simple illustration, a conventional (classical) computer uses machine learning to classify images of animals. Images whose pixels 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 team[3] 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

natural bridge between machine learning and quantum theory. However, recognizing this bridge is only the beginning.

For instance, it remains to be seen whether the way in which Havlíček *et al.* represent data in quantum space is actually useful for real-world machine-learning applications. That is, it is not known whether the approach is associated with a meaningful measure of similarity that, for example, in classifying images of animals, places cat pictures close to cat pictures but not to dog pictures. Moreover, it is unclear whether there are other strategies that would work better. And would these techniques be good enough to beat almost 30 years of classical methods? If so, the desperate search for a 'killer application' for quantum computers would be over. But the answer to this question is probably more complicated. ∎

Maria Schuld *is at Xanadu Quantum Technologies, Toronto, Ontario M5V 2L7,*

Canada, and at the Quantum Research Group, School of Chemistry and Physics, University of KwaZulu-Natal, Durban, South Africa.
e-mail: maria@xanadu.ai

1. Havlíček, V. *et al. Nature* **567**, 209–212 (2019).
2. Schölkopf, B. & Smola, A. J. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond* (MIT Press, 2002).
3. Schuld, M. & Killoran, N. *Phys. Rev. Lett.* **122**, 040504 (2019).

MEDICAL RESEARCH

Molecular envoys aid cancer spread

Pancreatic cancer usually spreads to the liver. The identification of signals from cells adjacent to pancreatic tumours that boost liver colonization might suggest ways to block this deadly form of cancer invasion. SEE LETTER P.249

ANIRBAN MAITRA

Pancreatic cancer is rapidly lethal, and the five-year post-diagnosis survival rate in the United States is 8% (ref. 1). At diagnosis, the cancer has usually already spread beyond its primary pancreatic site to invade other parts of the body, most commonly the liver[2]. This renders futile the option of surgically removing the pancreatic tumour to prevent such lethal spread, or metastasis[3]. On page 249, Lee *et al.*[4] report their identification, in mice and humans, of molecules made in the pancreas that travel to the liver and alter its environment to create conditions that assist cancer-cell invasion.

Much remains to be uncovered about the signals and sequence of events that precede and facilitate establishment of the implantation site for tumour invasion — known as the pro-metastatic niche[5]. Alterations that enable niche formation include blood-vessel changes that create cancer-cell docking sites and modifications to the layer of endothelial cells that form an outer barrier around tissues and that must be crossed for tissue invasion[5].

Although metastasis is usually the main reason for the failure of cancer treatment and for eventual death, it is a remarkably inefficient process. Cancers release millions of cells into the bloodstream each day, yet studies of skin cancer in animal models indicate that fewer than 0.1% of tumour cells form metastases[6]. For metastasis to be successful, cancer cells must exit their primary site, enter the bloodstream and overcome challenges that include surviving physical stress in blood vessels, adapting to the unfamiliar cellular surroundings of a different host organ, and evading destruction by immune cells. Therefore, understanding the factors that create a pro-metastatic niche are of crucial importance for clarifying how cancer cells overcome such obstacles to become established at a distant site.

Lee and colleagues investigated how pancreatic-tumour cells generate the pro-metastatic niche. The authors demonstrate that, in mice, the protein interleukin 6 (IL-6), a type of immune-signalling molecule called a cytokine, is secreted from non-cancerous fibroblast cells[7] in the microenvironment of the pancreatic tumour cells (Fig. 1). Fibroblasts are the main cells of the connective tissue. The authors report that IL-6 binds to its receptor protein on liver cells and drives expression of the transcription-factor protein STAT3, which is then activated by undergoing phosphorylation (the addition of a phosphate group). Liver cells that express such activated STAT3 secrete the proteins SAA1 and SAA2, which prepare the liver for the influx of cancer cells. The SAA proteins attract myeloid cells, which dampen the body's immune-surveillance response by secreting cytokines that inhibit cancer-killing T cells. SAA1 and SAA2 also drive the activation of hepatic stellate cells, a type of liver cell that deposits extracellular-matrix material, thereby aiding the initial anchoring and
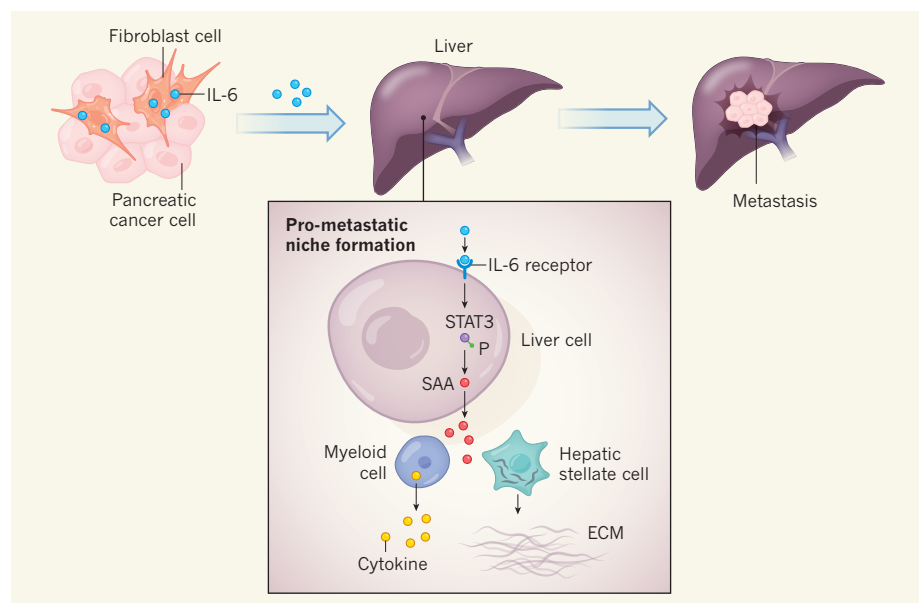
Figure 1 | A signal from the pancreas aids cancer invasion of the liver. Lee *et al.*[4] report studies in mice and humans that have uncovered a process driving the deadly step of cancer spread. The authors report that the protein IL-6, which is synthesized in non-cancerous fibroblast cells adjacent to a pancreatic cancer, is a key driver of tumour invasion of the liver. IL-6 travels through the bloodstream to the liver, where it binds to its receptor on liver cells. This drives expression of the protein STAT3, which is then phosphorylated (P denotes a phosphate group), and triggers the expression of SAA proteins (SAA1 and SAA2). These proteins are secreted from the cell and attract myeloid cells, which express cytokine molecules that dampen immune responses. SAA proteins also activate hepatic stellate cells, which deposit extracellular-matrix material (ECM). These changes create an environment, termed a pro-metastatic niche, that supports cancer colonization and growth. Once the pro-metastatic niche has formed, pancreatic cancer cells can invade the liver to form a secondary tumour site (metastasis).