back and forth between them takes time and power, and creates a bottleneck in performance.

To take advantage of AI technology, hardware engineers are looking to build computers that go beyond the constraints of von Neumann design. This would be a big step forward. For decades, advances in computing have been driven by scaling down the size of the components, guided by Gordon Moore's prediction that the number of transistors on a chip roughly doubles every two years — which generally meant that processing power did the same.

Modern computers bear little resemblance to early machines that used punch cards to store information and mechanical relays to perform calculations. Integrated circuits now contain transistors so small that more than 100 million of them would fit on the head of a pin. Yet the fundamental design of separate memory and processing remains, and that places a limit on what can be achieved.

One solution could be to merge the memory and processing units, but performing computational tasks within a memory unit is a major technical challenge.

Google's AlphaGo research shows a possible, different, way forward. The company has produced new hardware called a tensor processing unit, with an architecture that enables many more operations to be performed simultaneously. This approach to parallel processing significantly increases the speed and energy efficiency of computationally intensive calculations. And designs that relax the strict need to perform exact and error-free computation — a change in strategy known as approximate computing — could increase these benefits further.

As a result, the power consumption of AI programs such as AlphaGo has improved dramatically. But increasing the energy efficiency of such hardware is essential for AI to become widely accessible.

The human brain is the most energy-efficient processor around, so it is natural for hardware developers to try to mimic it. An approach called neuromorphic computing aims to do just that, with technologies that seek to simulate communication and processing in a biological nervous system. Several neuromorphic systems have already demonstrated the ability to emulate collections of neurons on tasks such as pattern recognition.

These are baby steps, and now the SRC has stepped in to try to encourage the hardware to walk. Under its Joint University Microelectronics Program, the SRC has quietly placed its focus on developing hardware architecture. A new centre at Purdue University in West Lafayette, Indiana, for example will research neuromorphic computing, and one at the University of Virginia in Charlottesville will develop ways of harnessing computer memory for extra processing power.

"The fundamental design of separate memory and processing places a limit on what can be achieved."

This technological task is huge. So it is heartening to see the SRC, traditionally US-centric, opening its doors. South Korean firm Samsung joined in late 2017, the fifth foreign company to sign up in the past two years. This is a welcome sign of collaboration. But that commercial rivals would work together in this way also signals how technically difficult the industry thinks it will be to develop new hardware systems.

As this research develops, *Nature* looks forward to covering progress and publishing results. We welcome papers that will enable computing architectures beyond von Neumann, such as components for neuromorphic chips and in-memory processing. Scientists across many fields are waiting for the result: computers powerful enough to sift all of their new-found data. They will have to wait a while yet. But the wait should be worth it. ■

Maths revision

A decadal update of academic mathematics shows the value of taking one's time.

Mathematics has its own way of doing things. Not for mathematicians the breakneck chase after the latest academic fad. "It goes up and down over the centuries," said one expert, when asked whether fluid dynamics — her focus — is now trendy.

Maths moves at its own pace, and the field is currently involved in a global effort to analyse, audit and agree new classifications of how mathematicians study and make use of maths. The MSC2020 system, due to appear in 2020, will formally approve new categories of maths, and split existing definitions into finer classes.

MSC stands for Mathematics Subject Classification, and it provides taxonomical order. In the current MSC2010, for instance, the code 03 represents mathematical logic and foundations. Going deeper, 03E is set theory and 03E72 is fuzzy set theory.

Why bother? The system is jointly managed by the mathematical resource zbMATH, curated by the Leibniz Institute for Information Infrastructure in Karlsruhe, Germany, and by the American Mathematical Society's *Mathematical Reviews*. Each is a 'meta-journal' that systematically summarizes and reviews every paper that comes out in the peer-reviewed mathematical literature. *Mathematical Reviews* and zbMATH use the MSC in their internal workflows, and many other journals have adopted the system to assign submissions to editors and reviewers. Mathematicians also use the numerical codes to search for papers in their speciality.

To keep the system up to date, every ten years the two organizations consult reviewers and request suggestions for new entries from the broader community. Nominations opened in July 2016 and close this August. A theme emerging for proposed new categories is for fields that mix traditional disciplines — such as 'algebraic statistics' and 'numerical algebraic geometry'.

Take topological data analysis, a popular candidate for inclusion. The theory has its roots in topology — the study of shapes and their arrangements within one another — which includes knot theory and higher-dimensional spaces. For more than a century, topology was mostly a pure-maths affair. But researchers have found ways to use it to give structure to large data sets, and so topological data analysis has been born.

More generally, the revision takes the pulse of broader cultural shifts. Suggested new categories indicate that more mathematicians have started to collaborate with researchers in other fields.

Recognition of a new subfield can depend on building citations, and that is a slow process in maths. A recent study of some 20 million references for more than 900,000 mathematical articles in zbMATH found that the time it takes for a paper's citations to peak is several years longer than in other fields — and is lengthening. Consequently, it takes a while for even the most dramatic breakthroughs to register in the MSC system. Many mathematicians expect Peter Scholze, a number theorist at the University of Bonn in Germany, to win a Fields Medal this year for his pioneering work on perfectoid spaces. But, as a research category, perfectoid spaces — only around since 2010 or so — is probably too undercooked yet to make the cut for MSC2020.

Can such a rigid hierarchy survive in an age of fluid metadata and keyword tagging? For now, it remains relevant. Studies have found a high correlation between clustering of the mathematical literature into topics — as measured from citation networks — and the MSC, at least at its upper levels. But things might change. For its own journals, for example, the American Physical Society changed in 2016 from a system similar to the MSC to a hybrid one called Physics Subject Headings. This has both a hierarchical tree of subfields and a broader set of 'facets' that cut across them like a Venn diagram, encompassing many terms. Maths might do the same at some point — but, quite correctly, in its own time. Maths has no need to start following fashion now. ■

CORRECTION

The Editorial 'Maths revision' (*Nature* **554**, 146; 2018) mistranslated the name of the Leibniz Institute. It is actually the Leibniz Institute for Information Infrastructure.