

populations have diverged from this microbial signature (Fig. 1).

The authors moved beyond focusing on species identity: they compared the genes, and the predicted functions of the proteins encoded by those genes, for the microbes in palaeofaeces with those found in present-day samples. Both the industrialized and the non-industrialized present-day samples had a greater prevalence of antibiotic-resistance genes than did the palaeofaeces, a finding consistent with the ancient microbes being from before the era of antibiotic use. Palaeofaeces had a high prevalence of genes encoding proteins that can degrade the molecule chitin, a component of insect exoskeletons. This finding is consistent with human consumption of insects, known to be a component of ancestral diets. Insect ingestion was confirmed by the authors' microscopy analysis of material in the palaeofaeces. The authors report many genes that were particularly prevalent in industrialized samples, including those involved in the degradation of mucus in the human gut.

Wibowo and colleagues' study is a remarkable technical achievement. They were able to recover high-quality DNA from microbial organisms that lived thousands of years ago, probably because of the good preservation possible in the dry desert environment in which the samples were located. Multiple independent lines of evidence authenticated the sample age and the human origin of the faeces. Having these ancient DNA sequences available in the public domain will undoubtedly benefit scientists for years to come.

However, DNA-sequence-based analyses do have limitations when the results are not paired with validation by other types of laboratory experiment. Using computational tools to predict information about proteins encoded by DNA is an imperfect method under ideal conditions, and is particularly tricky when analysing gene functions for previously unknown organisms, such as those discovered in this study. Moreover, microbiomes are highly variable between individuals and between populations. Analyses of more palaeofaeces from a wider range of timescales and locations will be needed to better understand general and population-specific features of ancient human gut microbiomes.

The authors found notable differences in the composition and function of microbes in palaeofaeces compared with those of microbes in present-day faeces. The higher prevalence of mucus-degrading species and genes in industrialized microbiomes than in ancient and non-industrialized ones is probably driven by Western diets, which often lack sufficient dietary fibre to support once-numerous fibre-degrading microbial species^{11,12}. Given the links between the microbiome and the immune system, these differences might be connected to the rising rates of

autoimmune, inflammatory and metabolic disorders in industrialized populations^{9,10}.

Wibowo and colleagues' work indicates that there are now two viable alternatives to time travel for understanding the composition of ancient microbiomes. Palaeofaeces enable the direct investigation of ancient microbiomes, but the sample age limits the further measurements and experiments that can be performed. Importantly, this study validates that present-day Indigenous populations living traditional lifestyles have similar microbiome compositions to those of ancient humans. It is essential to acknowledge that most of these present-day populations are marginalized, lead a vulnerable existence, and require exceptional protections to ensure they are not exploited. With ethically conducted research, these modern populations might open a window on our microbial past.

Matthew R. Olm and **Justin L. Sonnenburg** are in the Department of Microbiology and

Computer engineering

Superhuman floorplans for microchips

Andrew B. Kahng

A machine-learning system has been trained to place memory blocks in microchip designs. The system beats human experts at the task, and offers the promise of better, more-rapidly produced chip designs than are currently possible. **See p.207**

Success or failure in designing microchips depends heavily on steps known as floorplanning and placement. These steps determine where memory and logic elements are located on a chip. The locations, in turn, strongly affect whether the completed chip design can satisfy operational requirements such as processing speed and power efficiency. So far, the floorplanning task, in particular, has defied all attempts at automation. It is therefore performed iteratively and painstakingly, over weeks or months, by expert human engineers. But on page 207, researchers from Google (Mirhoseini *et al.*¹) report a machine-learning approach that achieves superior chip floorplanning in hours.

Modern chips are a miracle of technology and economics, with billions of transistors laid out and interconnected on a piece of silicon the size of a fingernail. Each chip can contain tens of millions of logic gates, called standard cells, along with thousands of memory blocks, known as macro blocks, or macros. The cells and macro blocks are interconnected by tens

of kilometres of wiring to achieve the designed functionality. Given this staggering complexity, the chip-design process itself is another miracle – in which the efforts of engineers, aided by specialized software tools, keep the complexity in check.

The locations of cells and macro blocks in the chip are crucial to the design outcome. Their placement determines the distances that wires must span, and thus affects whether the wiring can be successfully routed between components and how quickly signals can be transmitted between logic gates. Optimization of chip placement has been extensively studied for at least six decades^{2,3}. Seminal innovations in the mathematical field of applied optimization, such as a method known as simulated annealing⁴, have been motivated by the challenge of chip placement.

1. Hooper, L. V., Littman, D. R. & Macpherson, A. J. *Science* **336**, 1268–1273 (2012).
2. Karlsson, F., Tremaroli, V., Nielsen, J. & Bäckhed, F. *Diabetes* **62**, 3341–3349 (2013).
3. Asnicar, F. *et al.* *mSystems* **2**, e00164-16 (2017).
4. Moeller, A. H. *et al.* *Science* **353**, 380–382 (2016).
5. Vangay, P. *et al.* *Cell* **175**, 962–972 (2018).
6. Dethlefsen, L. & Relman, D. A. *Proc. Natl Acad. Sci. USA* **108**, 4554–4561 (2011).
7. Wibowo, M. C. *et al.* *Nature* **594**, 234–239 (2021).
8. Smits, S. A. *et al.* *Science* **357**, 802–806 (2017).
9. Blaser, M. J. *Cell* **172**, 1173–1177 (2018).
10. Sonnenburg, J. L. & Sonnenburg, E. D. *Science* **366**, eaaw9255 (2019).
11. Makki, K., Deehan, E. C., Walter, J. & Bäckhed, F. *Cell Host Microbe* **23**, 705–715 (2018).
12. Desai, M. S. *et al.* *Cell* **167**, 1339–1353 (2016).

The authors declare no competing interests. This article was published online on 17 May 2021.

Because macro blocks can be thousands or even millions of times larger than standard cells, placing cells and blocks simultaneously is extremely challenging. Modern chip-design methods therefore place the macro blocks




Figure 1 | Microchip floorplans designed by humans differ from those produced by a machine-learning system. An early step in microchip design is floorplanning – the placement of memory components called macro blocks on an empty layout canvas. Floorplanning is immensely complicated because of the vast number of potential configurations of macro blocks, and it involves multiple iterations as the logic-circuit design evolves. Each iteration is produced manually by human engineers, over days or weeks. **a**, This floorplan for a chip (the Ariane RISC-V processor⁸) is considered by human designers to be a good one. Its 37 macro blocks are close-packed in well-aligned rows and columns, leaving an uncluttered area for placement of other components. **b**, Mirhoseini *et al.*¹ report a machine-learning agent that, in just a few hours, designs floorplans that outperform those designed by humans. This agent-produced arrangement is another implementation of the Ariane processor, and is very different from that shown in **a**. (Image in **b** adapted from Extended Data Fig. 4 of the paper¹.)

first, in a step called floorplanning. Standard cells are then placed in the remaining layout area. Just placing the macro blocks is incredibly complicated: Mirhoseini *et al.* estimate that the number of possible configurations (the state space) of macro blocks in the floorplanning problems solved in their study is about $10^{2,500}$. By comparison, the state space of the black and white stones used in the board game Go is just 10^{360} .

Viable floorplanning solutions must leave empty regions on the chip to achieve all of the subsequent steps – placement of the standard cells, routing of the wiring and maximizing of the chip's processing speed. However, the optimizations of logic circuitry inherent in these steps can increase the total area taken up by standard cells by 15% or more. Human engineers must therefore iteratively adjust their macro-block placements as the logic-circuit design evolves. Each of these iterations is carried out manually, and takes days or weeks.

The computer industry has famously been driven by Moore's law – the number of components per chip has roughly doubled every two years. This rate of advancement corresponds to an increase in the number of components on a chip of about one per cent per week. The failure to automate floorplanning is therefore problematic – not only because of the associated time costs, but also because it limits the number of solutions that can be explored within chip-development schedules.

But everything changed on 22 April 2020. On that day, Mirhoseini *et al.* posted a preprint⁵

of the current paper to the online arXiv repository. It stated that “in under 6 hours, our method can generate placements that are superhuman or comparable” – that is, the method can outperform humans in a startlingly short period of time. Within days, numerous semiconductor-design companies, design-tool vendors and academic-research groups had launched efforts to understand and replicate the results.

Mirhoseini and colleagues trained a machine-learning ‘agent’ that can successfully place macro blocks, one by one, into a chip layout. This agent has a brain-inspired architecture known as a deep neural network, and is trained using a paradigm called reinforcement learning. At any given step of floorplanning, the trained agent assesses the ‘state’ of the chip being developed, including the partial floorplan that it has constructed so far, and then applies its learnt strategy to identify the best ‘action’ – that is, where to place the next macro block.

The technical details of this approach, such as how to represent the chip-design and partial-floorplanning solutions, were developed with the overarching goal of finding a general, transferable solution to the macro-placement problem. In other words, the trained agent should succeed even when confronted with chip designs that it has not previously encountered, drawn from a wide range of applications and markets. The authors report that, when their agent is pre-trained on a set of 10,000 chip floorplans, it is already quite successful when used in a ‘one shot’ mode on a new

design: with no more than six extra hours of fine-tuning steps, the agent can produce floorplans that are superior to those developed by human experts for existing chips. Moreover, the agent's solutions are very different from those of trained human experts (Fig. 1).

Arthur C. Clarke famously noted⁶ that “any sufficiently advanced technology is indistinguishable from magic”. To long-time practitioners in the fields of chip design and design automation, Mirhoseini and colleagues' results can indeed seem magical. In the past year, experts worldwide have contemplated questions such as, ‘How is it that the agent can initially place each macro block in turn so effectively that the chosen placement is used in the final, manufactured chip design?’

The authors report that the agent places macro blocks sequentially, in decreasing order of size – which means that a block can be placed next even if it has no connections (physical or functional) to previously placed blocks. When blocks have the same size, the agent's choice of the next block echoes the choices made by ‘cluster-growth’ methods⁷, which were previously developed in efforts to automate floorplan design, but were abandoned several decades ago. It will be fascinating to see whether the authors' use of massive computation and deep learning reveal that chip designers took a wrong turn in giving up on sequential and cluster-growth methods.

Another much-debated question has been, ‘How does the agent's choice of macro-block placements survive subsequent steps in the chip-design process?’ As mentioned earlier, human engineers must iteratively adjust their floorplans as the logic-circuit design evolves. The trained agent's macro-block placements somehow evade such landmines in the design process, achieving superhuman outcomes for timing (ensuring that signals produced in the chip arrive at their destinations on time) and for the feasibility and efficiency with which wiring can be routed between components. Moreover, Mirhoseini and colleagues' use of simple metrics as proxies for key parameters of the chip design works surprisingly well – it will be interesting to understand why these proxies are so successful. The authors' intention to make their code available is invaluable in this light.

The development of methods for automated chip design that are better, faster and cheaper than current approaches will help to keep alive the ‘Moore's law’ trajectory of chip technology. Indeed, for technical leaders and decision-makers in the chip industry, the most important revelation in Mirhoseini and colleagues' paper might be that the authors' floorplan solutions have been incorporated into the chip designs for Google's next-generation artificial-intelligence processors. This means that the solutions are good enough for

millions of copies to be printed on expensive, cutting-edge silicon wafers. We can therefore expect the semiconductor industry to redouble its interest in replicating the authors' work, and to pursue a host of similar applications throughout the chip-design process.

Andrew B. Kahng is in the Department of Computer Science and Engineering, and in the Department of Electrical and Computer Engineering, University of California, San Diego, La Jolla, California 92093, USA.
e-mail: abk@ucsd.edu

1. Mirhoseini, A. et al. *Nature* **594**, 207–212 (2021).
2. Markov, I. L., Hu, J. & Kim, M.-C. *Proc. IEEE* **103**, 1985–2003 (2015).
3. Kahng, A. B. *Proc. ACM/IEEE Int. Symp. Physical Design* 15–22 (2021).
4. Kirkpatrick, S., Gelatt, C. D. Jr & Vecchi, M. P. *Science* **220**, 671–680 (1983).
5. Mirhoseini, A. et al. Preprint at <https://arxiv.org/abs/2004.10746> (2020).
6. Clarke, A. C. *Profiles of the Future: An Inquiry into the Limits of the Possible* 41 (Macmillan, 1973).
7. Sait, S. M. & Youssef, H. *VLSI Physical Design Automation: Theory and Practice* Ch. 3, 90–95 (McGraw-Hill, 1995).
8. Zaruba, F. & Benini, L. *IEEE Trans. Very Large Scale Integration* **27**, 2629–2640 (2019).

The author declares no competing interests.

Evolution

Seed for thought

Douglas E. Soltis

The origin and rapid diversification of flowering plants is a long-standing “abominable mystery”, as Charles Darwin put it. Part of the puzzle – the origin of the protective covering of flowering-plant seeds – is nearing resolution. **See p.223**

One of the landmark events in evolution is the appearance of flowering plants, termed angiosperms. On page 223, Shi *et al.*¹ describe fossil evidence that sheds light on a long-standing mystery about seed evolution.

A number of botanical innovations accompanied the appearance of flowering plants and contributed to the organisms' rapid rise to worldwide dominance of terrestrial and many aquatic ecosystems. Such innovations include floral organs and a special nutritive tissue for the plant embryo – endosperm. In addition, flowering plants and another plant group, gymnosperms, produce seeds, which provide a protective layer around the developing embryo. This enabled seed plants to truly conquer the terrestrial environment, overtaking other land plants such as mosses and ferns, which do not produce seeds. Seeds have also had a central role in enabling human survival by providing an important food source.

Flowering plants arose from an ancestor in the gymnosperms^{2,3}. The fossil record contains many clades of now-extinct gymnosperms, and it is not clear which group gave rise to the angiosperms, although it is definitely not one of the living groups of gymnosperms – plants such as conifers, ginkgo (*Ginkgo biloba*) or cycads^{2,3}.

All gymnosperms have just one protective layer, termed an integument, that surrounds their seeds, whereas flowering plants have two such layers. In a letter⁴ to the botanist J. D. Hooker, Charles Darwin described the origin of flowering plants as an “abominable mystery”. An enigma embedded in this

mystery is how the second (outer) protective layer evolved. The outer integument differs from the inner one in its developmental pathway and in the genetic control governing tissue-layer formation^{5,6}, and the two

integuments are thus clearly very different from each other.

Shi and colleagues describe extremely well-preserved fossils of extinct seed plants. They use these remarkable newly described fossils, as well as other similar-looking fossils reported earlier², to build a tree of relationships between plants. Their effort places all of these fossils on the plant family tree at a position close to that of flowering plants – they are very close relatives of modern angiosperms, and, as such, might provide clues to their origin. The fossils of these ancient relatives of flowering plants show remarkable diversity in the shape of their reproductive structures. Amazingly, these are not recently discovered fossil specimens; they were collected almost a century ago, deposited in museum collections and only recently unearthed for a second time (this time, from museum drawers by the authors), with their relevance for reconstructing the plant family tree now finally recognized.

These extinct plants had a cup-like structure, termed a cupule (Fig. 1), that surrounded the developing seed or seeds (each seed itself having just one integument, which typifies all gymnosperms, living and extinct). The role of these cupules is unknown, but they might have provided extra protection for the seed, or aided its dispersal.

Might the cupule be the precursor to the outer protective layer unique to flowering-plant seeds? There is another twist (pun intended) to solving the mystery of angiosperm seeds.

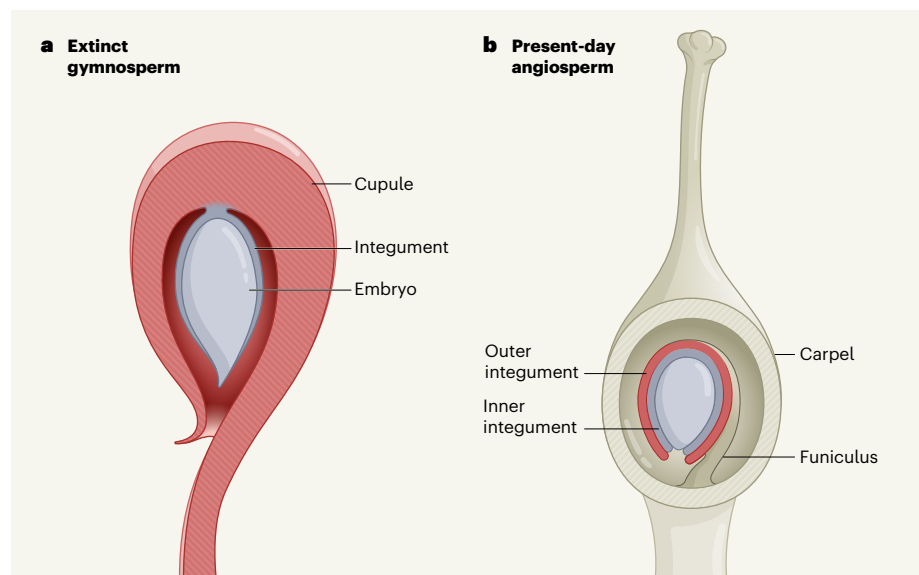


Figure 1 | Seed evolution. Shi *et al.*¹ present fossil evidence that sheds light on a long-standing mystery about the origin of the protective layer, termed the outer integument (or second integument), that surrounds the seeds of flowering plants (angiosperms). It is thought that flowering plants arose from now-extinct plants belonging to a group called the gymnosperms (living members of which include conifers). **a**, Gymnosperms have one integument that surrounds the embryo of a seed. A structure called the cupule, which might have provided protection or aided seed dispersal, formed the outer layer of ancient gymnosperm seeds. **b**, Flowering plants have inner and outer integuments. Shi and colleagues find that the cupule probably evolved to form the outer integument of flowering plants. Angiosperm seeds are connected to the carpel structure that surrounds them through a stalk called the funiculus. Shi and colleagues suggest that the funiculus evolved from the stalk of the gymnosperm cupule.