memory technology itself, the circuits that connect various chip elements and the architecture — the way in which the entire chip is laid out. Most research on this topic still focuses only on these steps. But the technology also requires the development of a compiler that can translate code into machine-level instructions, algorithms that are distinct from those developed for digital chips, and applications that are optimized for analog chips.

Even at the level of memory technology, the best way forwards is unclear. Phase-change memories are not the only options for analog chips. They are attractive in that they are non-volatile, which means that information can be retained even after the power is switched off. However, this feature can also be a burden. If the weights need to be reconfigured — for example, when the model size exceeds the on-chip memory or when AI training necessitates weights to be refreshed frequently — then non-volatile devices consume more power than do volatile memories, because the task needs higher voltages and complicated programming procedures with non-volatile devices. Volatile memories have therefore also been considered for use in analog chips, and have demonstrated higher efficiency than have non-volatile memories in some cases[3].

Circuit innovation is the next step to consider. A flaw in most analog-AI implementations so far is that they focus only on the multiply–accumulate operation and leave all other computing tasks in the digital domain. This means that data need to be converted from analog to digital, and vice versa, which slows computing down and limits performance. To overcome this, researchers either need to invent new techniques for converting data[4] or bring more digital operations into the analog domain.

After the circuit, the challenge is getting the architecture right. In the early 2010s, it became clear that GPUs were more efficient than CPUs for some applications. Analog AI chips represent the next step in this evolution: their throughput and efficiency are considerably better than those of CPUs and GPUs, but this comes at the expense of flexibility. A hybrid analog–digital architecture is the remedy, because digital components' flexibility enables them to bridge gaps that analog devices cannot, such as assigning computational resources[5] and correcting computation or storage errors.

These three steps would provide the hardware foundation for an analog AI chip. To further improve Ambrogio and colleagues' chip and unleash its full potential, the remaining three steps must also be tackled. The astonishing efficiency of the authors' chip merely reflects a theoretical maximum — in practice, the percentage of analog AI hardware that is actually used during computations can be very limited[6]. A customized compiler is therefore essential because it can segment tasks, then map each task efficiently to the available hardware to maximize performance.

Tailored algorithms are similarly crucial. Analog computing is inherently prone to generating errors because it is vulnerable to problems such as thermal noise, manufacturing imperfections and variations in the thermal and electrical environment of the device. This means that performance is compromised when analog-AI chips use algorithms that are designed for conventional digital computing — an issue for which there are two promising solutions. First, researchers can mitigate the impact of errors by using algorithm optimization techniques to relax the required computational precision[7]. Alternatively, they can embrace algorithms that leverage analog errors, such as are involved in Bayesian neural networks[8], which use statistical inference methods to improve the performance of ordinary neural networks.

Finally, developing dedicated applications for analog-AI chips is a key step towards making them commercially viable — and it is a challenging one. It took decades to shape the computational ecosystems in which CPUs and GPUs operate so successfully, and it will probably take years to establish the same sort of environment for analog AI. The good news is that Ambrogio and colleagues, together with other researchers in this area, are steering the ship, and have set sail towards realizing this goal.

Hechen Wang is at Intel Labs, Intel Corporation, Hillsboro, Oregon 97124, USA.
e-mail: hechen.wang@intel.com

1. Ambrogio, S. *et al*. *Nature* **620**, 768–775 (2023).
2. Horowitz, M. *IEEE Int. Solid-State Circuits Conf. Digest Tech. Pap.* **57**, 10–14 (2014).
3. Wang, H. *et al*. *IEEE J. Solid-State Circuits* **58**, 1037–1050 (2023).
4. Wu, P.-C. *et al*. *IEEE Int. Solid-State Circuits Conf. Digest Tech. Pap.* **66**, 126–128 (2023).
5. Chen, G. K., Knag, P. C., Tokunaga, C. & Krishnamurthy, R. K. *IEEE J. Solid-State Circuits* **58**, 1117–1128 (2023).
6. Sze, V., Chen, Y.-H., Yang, T.-J. & Emer, J. S. *IEEE Solid-State Circuits Mag.* **12**, 28–41 (2020).
7. Keller, B. *et al*. *IEEE J. Solid-State Circuits* **58**, 1129–1141 (2023).
8. Dorrance, R., Dasalukunte, D., Wang, H., Liu, R. & Carlton, B. *IEEE J. Solid-State Circuits* https://doi.org/10.1109/JSSC.2023.3283186 (2023).

The author declares no competing interests.

Genetics

Nuclear genome regulates mitochondrial DNA

Sonia Boscenco & Ed Reznik

A genetic analysis provides the most-detailed glimpse yet of how genetic variants in nuclear DNA regulate the copy number and variability of DNA housed in organelles called mitochondria. **See p.839**

Mitochondria are cellular organelles that house their own DNA. Each DNA molecule is built from 16,569 nucleotides — a stark contrast to the 3 billion nucleotides needed to build each copy of the nuclear genome (nuDNA)[1]. There are thousands of copies of mitochondrial DNA (mtDNA) in each cell, and each copy encodes 13 proteins that are essential for cellular energy production[1]. The genetic integrity of mtDNA is therefore crucial for healthy cell function. However, the proteins required for maintaining and replicating mtDNA are encoded by nuDNA[2]. On page 839, Gupta *et al*.[3] shed light on how nuDNA mediates maintenance and regulation of mtDNA, with consequences for human disease, physiology and evolution.

Gupta and colleagues initially focused on a distinguishing aspect of mtDNA: its high copy number per cell. Shifts in mtDNA copy number (mtCN) are associated with ageing-related disorders[4] and cancer[5], and low mtCN in blood cells has been linked with an elevated risk of a range of diseases, from type 2 diabetes[6] to heart attack[7]. The authors searched for factors associated with mtCN variation using the UK Biobank — a large-scale publicly available genetic database. The group discovered that nearly 25% of the variation in mtCN between people in this cohort could be ascribed to blood-cell composition, potentially reflecting different mtCN levels in different blood-cell types. After adjusting for blood-cell composition and other covariates, the associations between mtCN and disease that had previously been reported were no longer apparent. Instead, the authors found that mtCN simply decreased with age.

Next, Gupta *et al.* performed a genome-wide association study (GWAS) to identify nuDNA variants associated with variation in covariate-corrected mtCN. This analysis revealed previously undetected nuDNA variants associated with mtCN (for example, located in the *MCAT* gene, which encodes a protein implicated in mitochondrial metabolism), and validated existing ones (such as variants in the nuDNA-encoded gene *TFAM*, which encodes a protein essential for mtDNA replication).

A natural consequence of the high copy number of mtDNA is that, in any given cell, only some of the mtDNA copies might harbour a mutation of interest. This phenomenon, known as heteroplasmy, is of clinical interest because heteroplasmy of harmful mtDNA mutations is widely thought to affect the manifestation of disease[8]. The authors asked whether heteroplasmy might explain why, although nearly one in 200 individuals carry disease-causing mtDNA variants[9], only one in 5,000 exhibit the associated disease[10].

They studied the clinical profiles of people who carried well-described, harmful mtDNA variants, including the mutation at nucleotide position 3243 (referred to as m.3243A>G), which can cause MELAS syndrome – a devastating mitochondrial disease that affects the nerves and muscles. They found that individuals who carried the harmful m.3243 mutation at low heteroplasmies (meaning that a low proportion of mtDNA copies carry the mutation in a single cell) displayed intermediate disease traits, such as impaired vision and hearing. These findings are consistent with a dosage-sensitive mode of action for these variants, and emphasize that low-heteroplasmy mtDNA variants might cause intermediate disease.

Next, the authors investigated how and when mtDNA variants are inherited, by evaluating how heteroplasmy levels relate between sibling–sibling and parent–offspring pairs across two large cohorts (from the UK Biobank and a US repository called All of Us), totalling more than 250,000 individuals. This analysis revealed two characteristic and distinct ways in which heteroplasmic mtDNA variants arise. Single-nucleotide variants were predominantly acquired after birth and accumulated with age over time – especially after the age of 70. By contrast, insertions and deletions of mtDNA were chiefly inherited maternally and shared across siblings, and did not generally accumulate with age.

Gupta *et al.* then performed another GWAS, to look for associations between nuDNA variants and heteroplasmies. They identified numerous nuDNA variants associated with heteroplasmic dosage, including genes involved in nucleotide metabolism and mtDNA replication and maintenance. Interestingly, they found few genes that were significantly associated with the presence or absence of a heteroplasmy (rather than its dosage). This argues against the possibility that these variants associate with heteroplasmies because they cause elevated rates of mutation in mtDNA. Instead, the authors propose that these variants make it more likely that maternally inherited mtDNA molecules of particular lengths will be preserved through replication.

The most common heteroplasmy found in the cohorts was at position m.302, which acquires insertions and deletions of various lengths, producing a 'length heteroplasmy'. Intriguingly, this region of mtDNA is known to have a role in controlling the rates of mtDNA replication and transcription[11] – and, in line with this, Gupta and colleagues found that heteroplasmy at m.302 was strongly associated with mtCN. These observations point to a two-pronged model for mtCN regulation (Fig. 1), first, by nuDNA, through variation in genes associated with mtDNA maintenance and replication (known as regulation in *trans*), and second, by mtDNA, through heteroplasmic variation at positions such as m.302 (regulation in *cis*).

Finally, the authors analysed the mtDNA composition of 171 single cells from one person, and found that each cell harboured a unique mixture of mtDNA molecules, specifically with respect to the length heteroplasmy at m.302. This finding adds to the now overwhelming evidence that individual cells are reservoirs of mtDNA diversity[12], and emphasizes the breadth and complexity of mtDNA regulation by the nuclear genome.

Gupta and colleagues' findings have implications for our understanding of how mtDNA evolves. At the scale of human populations, variation in mtCN and heteroplasmy is controlled in large part by the genetic variation in the nuclear genome. But at the most granular resolution of a single cell from a single individual, mtDNA molecules differing by a single nucleotide probably attain distinct, steady-state copy numbers, which might affect their capacity to carry out the most basic of metabolic and signalling functions. New theories of evolutionary selection and dynamics that span these disparate regimes hold the promise of transforming our understanding of mitochondrial genetics and the complex biology that emerges from it.

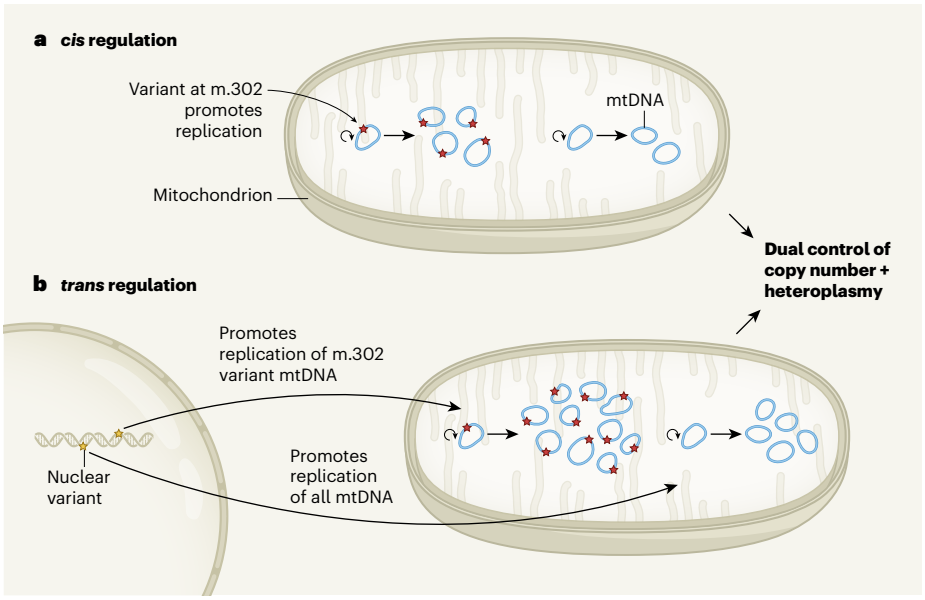

Figure 1 | Dual control of DNA copy number in organelles called mitochondria. a, Gupta *et al.*[3] report that maternally inherited variants in mitochondrial DNA (mtDNA) can affect replication of that DNA (known as *cis* regulation). For instance, certain variants at position m.302 are associated with an elevated number of mtDNA copies, perhaps because mtDNA molecules carrying these variants are more likely than others to be replicated. **b**, Variants in nuclear DNA can also affect mtDNA copy number (*trans* regulation), perhaps by selectively promoting the replication of certain m.302 variants, or by altering the behaviour of genes that encode proteins involved in mtDNA replication or maintenance in general. Together, this two-pronged regulatory system determines mtDNA copy number and the level of heteroplasmy (the degree to which different mtDNA sequences are present in a cell).

Sonia Boscenco and **Ed Reznik** are in the Department of Physiology, Biophysics and Systems Biology, Weill Cornell Medical College, New York, New York 10065, USA, and in the Computational Oncology Service, Memorial Sloan Kettering Cancer Center, New York.
e-mails: boscens@mskcc.org; reznike@mskcc.org

1. Stewart, J. B. & Chinnery, P. F. *Nature Rev. Genet.* **22**, 106–118 (2021).
2. Johnston, I. G. & Williams, B. P. *Cell Syst.* **2**, 101–111 (2016).
3. Gupta, R. *et al. Nature* **620**, 839–848 (2023).
4. Mengel-From, J. *et al. Hum. Genet.* **133**, 1149–1159 (2014).

5. Reznik, E. *et al. eLife* **5**, e10769 (2016).
6. Fazzini, F. *et al. J. Intern. Med.* **290**, 190–202 (2021).
7. Ashar, F. N. *et al. JAMA Cardiol.* **2**, 1247–1255 (2017).
8. Rossignol, R. *et al. Biochem. J.* **370**, 751–762 (2003).
9. Elliott, H. R., Samuels, D. C., Eden, J. A., Relton, C. L. & Chinnery, P. F. *Am. J. Hum. Genet.* **83**, 254–260 (2008).
10. Gorman, G. S. *et al. Ann. Neurol.* **77**, 753–759 (2015).
11. Wanrooij, P. H. *et al. Nucleic Acids Res.* **40**, 10334–10344 (2012).
12. Ludwig, L. S. *et al. Cell* **176**, 1325–1339 (2019).

E.R. declares competing interests. See go.nature.com/3kjkegk for details.
This article was published online on 16 August 2023.

Palaeontology

A really big fossil whale

J. G. M. Thewissen & David A. Waugh

A newly discovered fossil of an extinct whale from Peru indicates that the animal's skeleton was unexpectedly enormous. This finding challenges our understanding of body-size evolution. **See p.824**

Whales, dolphins and porpoises belong to a group called cetaceans, which includes the largest known animals to have ever lived on Earth. Until now it had been assumed that the blue whale (*Balaenoptera musculus*) holds the record for the largest body size. However, on page 824, Bianucci *et al.*[1] challenge that assumption in their presentation of fossil evidence from a 39-million-year-old whale found in Peru. This whale, a member of the basilosaurid group (a family of extinct cetaceans), not only had an extremely large body size, but also had an exceptionally heavy skeleton relative to its body mass. Discoveries of such extreme body forms are an opportunity to re-evaluate our understanding of animal evolution — it seems that we are only dimly aware of how astonishing whale form and function can be.

Vertebrate palaeontologists by necessity focus most of their attention on bone, and the bones of this newly discovered species termed *Perucetus colossus* are extremely unusual. The cross-section of a mammalian bone commonly looks like a baguette in terms of having a hard and solid crust (compact bone) that surrounds a spongy interior (trabecular bone). The proportions of crust and interior vary according to the animal's need. Hippopotamuses, for example, need to walk fully submerged on riverbeds, so their bones have a lot of compact bone that makes their skeletons heavy. Bones that are made up of a lot of compact bone are also common in slow-moving marine mammals such as manatees and bowhead whales (*Balaena mysticetus*).

But most whales are different — they are predators of fast-moving prey and have relatively light skeletons compared with other mammals of the same size (see Fig. 4 of ref. 1). *Perucetus colossus* is on the hippo end of the scale, with bones mostly or solely made of compact bone. This suggests it was not chasing fast prey, and the authors propose that it led an alternative lifestyle, that of a scavenger.

The discovery of *P. colossus* also invites us to think about how life-history strategies evolved. As the authors point out, gigantism characterizes whales that are active swimmers in the open ocean, but *P. colossus* was an exception to that as an extremely large whale that was unlikely to be a fast swimmer. Vertebrates that live underwater but are not bottom dwellers are usually more or less neutrally buoyant, which means that the weight of heavy tissues (such as bones) is balanced by that of light tissues (such as fat), although there can be wide variation in absolute mass between species and between individuals of the same species.

Although *P. colossus*' skeleton is incomplete, consisting of several vertebrae, ribs and pelvic bones, it had a lot of heavy bone and therefore it must have had a lot of lighter tissues, too. This is a fundamental difference compared with animals living on land for which all tissues contribute to weight that needs to be supported by body parts such as limbs. By contrast, in water, heavier tissues can be offset by lighter tissues to acquire neutral buoyancy, and total mass is less important.

What type of life-history strategy might *P. colossus* have had? Cetaceans, unlike most mammals, keep growing long past

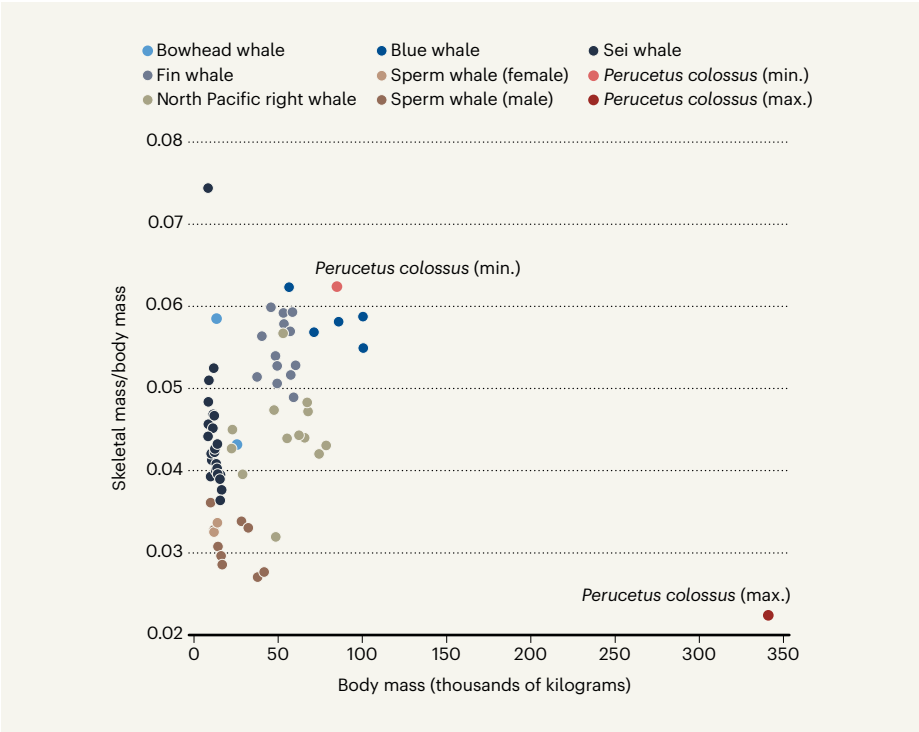

Figure 1 | The body mass and skeletal mass of individual whales. Whales grow for most of their lives, therefore individuals of a given species with a heavier body mass can typically be assumed to be older and for sperm whales (*Physeter macrocephalus*), males are heavier than the females. Bianucci *et al.*[1] discovered fossils of an extinct whale species that they named *Perucetus colossus*, and estimated its body mass and skeletal mass (minimum and maximum estimates are plotted). *Perucetus colossus* had a large body mass and an exceptionally heavy skeleton as revealed by the proportion of its total body mass that is attributed to bone (skeletal mass to body mass). Examining the proportion of body mass that is attributed to bone can shed light on a whale's lifestyle. Blue whales (*Balaenoptera musculus*) and fin whales (*Balaenoptera physalus*) are fast swimmers and published data[3–5] indicates that they maintain the same relative skeletal mass to body mass over their lifetime. By contrast, sei whales (*Balaenoptera borealis*)[5] and bowhead whales (*Balaena mysticetus*)[6–8] decrease their relative skeletal mass as they age. This is the type of lifestyle predicted for *P. colossus*. Bowheads, sei whales and North Pacific right whales (*Eubalaena japonica*)[9,10] might reduce blubber and bone in tandem as they age. Data shown are from refs 1, 3–10.