

Chicago, Illinois 60607, USA.
e-mail: lreyzin@uic.edu

1. Gödel, K. *Monatsh. Math.* **38**, 173–198 (1931).
2. Gödel, K. *The Consistency of the Continuum Hypothesis* (Princeton Univ. Press, 1940).
3. Cohen, P. J. *Proc. Natl Acad. Sci. USA* **50**, 1143–1148 (1963).
4. Cohen, P. J. *Proc. Natl Acad. Sci. USA* **51**, 105–110 (1964).
5. Ben-David, S., Hrubeš, P., Moran, S., Shpilka, A. & Yehudayoff, A. *Nature Mach. Intell.* <https://doi.org/10.1038/s42256-018-0002-3> (2019).
6. Viola, P. & Jones, M. J. *Int. J. Comput. Vis.* **57**, 137–154 (2004).
7. Valiant, L. G. *Commun. ACM* **27**, 1134–1142 (1984).
8. Vapnik, V. N. & Chervonenkis, A. Y. *Theory Probab. Appl.* **16**, 264–280 (1971).
9. Blumer, A., Ehrenfeucht, A., Haussler, D. & Warmuth, M. K. *Inf. Process. Lett.* **24**, 377–380 (1987).
10. Littlestone, N. & Warmuth, M. K. *Relating Data Compression and Learnability*, tech. rep. (Univ. California Santa Cruz, 1986).

models used in machine learning. In each case, the basic goal is to perform as well, or nearly as well, as the best predictor in a family of functions, such as neural networks or decision trees. For a given model and function family, if this goal can be achieved under some reasonable constraints, the family is said to be learnable in the model.

Machine-learning theorists are typically able to transform questions about the learnability of a particular function family into problems that involve analysing various notions of dimension that measure some aspect of the family's complexity. For example, the appropriate notion for analysing PAC learning is known as the Vapnik–Chervonenkis (VC) dimension⁸, and, in general, results relating learnability to complexity are sometimes referred to as Occam's-razor theorems⁹. These notions of dimension happen to be simple enough to leave no room for the spectre of unprovability to manifest itself. But Ben-David and colleagues show that machine learning cannot always escape this fate. They introduce a learning model called estimating the maximum (EMX), and go on to discover a family of functions whose learnability in EMX is unprovable in standard mathematics.

Ben-David *et al.* describe an example EMX problem: targeting advertisements at the most frequent visitors to a website when it is not known in advance which visitors will visit the site. The authors formalize EMX as a question about a learner's ability to find a function, from a given family, whose expected value over a target distribution is as large as possible. EMX is actually quite similar to the PAC model, but the slightly different learning criterion surprisingly connects it to the continuum hypothesis and brings unprovability into the picture.

The authors' proof involves a beautiful connection between machine learning and data compression that was first observed¹⁰ in the 1980s. The intuition is that, if a training sample labelled by a function from some family can always be compressed, the family must in some sense have low complexity, and therefore be learnable. Moreover, certain learning algorithms can be used to compress data. The authors introduce monotone compression — a variant of compression that they show to be appropriate for characterizing the learnability of particular function families in EMX.

Ben-David and colleagues then prove that the ability to carry out a weak form of monotone compression is related to the size of certain infinite sets. The set that the authors ultimately use in their work is the unit interval, which is the set of real numbers between 0 and 1. Their results imply that the finite subsets of the unit interval have monotone-compression schemes, and therefore are learnable in EMX, if and only if the continuum hypothesis is true, which is known to be unprovable.

Because EMX is a new model in machine learning, we do not yet know its usefulness for developing real-world algorithms. So these

results might not turn out to have practical importance. But we do now know that we should be careful when introducing new models of learning. Moreover, we might need to look again at the many subtleties that can come up, even in established learning models.

Machine learning has matured as a mathematical discipline and now joins the many subfields of mathematics that deal with the burden of unprovability and the unease that comes with it. Perhaps results such as this one will bring to the field of machine learning a healthy dose of humility, even as machine-learning algorithms continue to revolutionize the world around us. ■

Lev Reyzin is in the Department of Mathematics, Statistics, and Computer Science, University of Illinois at Chicago,

METABOLISM

Fat cells with a sweet tooth

Some fat cells convert energy into heat, so targeting them to induce weight loss is appealing. The discovery that a subset of the cells burns glucose, rather than both glucose and lipids, could improve our ability to do just that. [SEE ARTICLE P.180](#)

WENFEI SUN & CHRISTIAN WOLFRUM

Fat is often thought of as a means to store energy in the form of lipids. But this is just the role of white fat cells. The body also contains a second type of fat that burns the energy stored in nutrients to produce heat,

enabling mammals to maintain their body temperature in a cold environment¹. Activating this 'thermogenic' fat is thought to be an attractive way to combat obesity². On page 180, Chen *et al.*³ identify a previously uncharacterized type of thermogenic fat cell, which is derived from a hitherto unknown cell lineage




Figure 1 | Dual routes to heat-producing fat. **a**, White fat cells store energy. In wild-type mice under cold conditions, proteins called β -adrenergic receptors on fat-precursor cells (not shown) are activated by a β -adrenergic signalling pathway, which causes some precursors in white-fat tissues to differentiate into beige fat cells. This beige fat burns energy from both lipids and sugar to produce heat, maintaining body temperature. **b**, Chen *et al.*³ analysed β -less mice, which lack β -adrenergic receptors, meaning that β -adrenergic signalling is blocked. The authors showed that a different type of beige fat, dubbed glycolytic beige fat (g-beige fat), arises in these animals under cold conditions. The differentiation of these cells from a group of precursors that express the muscle-precursor protein MyoD (not shown) is driven by the transcription factor GABPa. The cells burn mainly glucose.

and burns mainly sugar, rather than the sugar and lipid combination processed by most thermogenic fat. Its discovery could open up new therapeutic possibilities for weight loss.

Thermogenic fat cells are currently classified as either brown or beige⁴. Unlike brown fat, beige fat cells are found interspersed in white fat. The production of beige fat cells can be triggered in response to cold by activation of the β -adrenergic-receptor proteins on the surface of precursor cells in white-fat tissue depots. These receptors are stimulated by the β -adrenergic signalling pathway, which originates in the nervous system.

Extensive efforts have been made to target thermogenic fat for weight loss, but no successful drug has been developed so far. This is due partly to the fact that the amount of active brown fat varies between individuals, and decreases with age. In addition, current strategies to induce the formation and activation of beige fat typically involve simulating cold responses, and are thought to act through β -adrenergic signalling. This inevitably affects other organs, because the β -adrenergic receptors are widely expressed across various tissues⁵ — and such a lack of specificity raises safety issues. Thus, other ways of inducing the formation and activation of thermogenic fat would be very valuable.

Chen *et al.* set out to identify alternative pathways by which to activate thermogenic fat using a mouse model called the β -less mouse, which lacks all three β -adrenergic receptors. They confirmed previous reports^{6,7} that some beige fat cells were still produced under these conditions (Fig. 1).

Where do these cells come from? Chen *et al.* performed an in-depth characterization of gene expression in the fat tissues of the β -less mice, which revealed that genes involved in muscle development were more highly expressed in fat tissues lacking β -adrenergic receptors than in those of control animals. The authors genetically engineered mice so that cells expressing one such gene, *Myod* (which encodes the protein MyoD and is normally expressed in muscle-cell precursors), were indelibly labelled with a fluorescent protein. They then blocked β -adrenergic signalling using a drug and exposed the mice to mild cold, before tracking the fate of MyoD-expressing cells and their descendants. This lineage-tracing experiment revealed that some MyoD-expressing cells give rise to a subset of beige fat cells, suggesting that muscle-cell precursors can be reprogrammed to turn into specific beige fat cells.

When they analysed the gene-expression profiles of the cells, Chen *et al.* found that this subset of beige fat differs from that of conventional beige fat. The subset expresses higher levels of many genes involved in sugar and carbohydrate metabolism and glycolysis — the process by which energy is produced from glucose. On the basis of this profile, the authors dubbed the cells glycolytic beige fat (g-beige

fat). The data indicated that the transcription factor GABPa drives the differentiation of MyoD-expressing progenitors into g-beige fat cells. The authors confirmed this supposition *in vitro*, showing that overexpression of GABPa in muscle progenitors leads to their differentiation into fat.

Chen *et al.* next demonstrated that g-beige fat has a physiological role in mice, mainly burning sugar to produce heat. G-beige fat was generated in wild-type animals subjected to prolonged, harsh cold conditions. Moreover, animals engineered so that they could not produce g-beige fat showed reduced glucose uptake and oxygen consumption in fat compared with controls; their ability to control their body temperature in response to cold was also impaired.

The authors' findings are exciting for several reasons. First, this study reinforces the idea that mature fat is composed of various cell types^{8,9}. It is plausible that other subpopulations of brown, beige and even white fat cells exist. These cell types might have different roles in regulating body-wide metabolism. Furthermore, it is possible that the g-beige fat cells have functions besides thermogenesis. For example, in recent years it has become evident that brown fat communicates with other tissues by secreting specific signalling molecules¹⁰ — perhaps g-beige fat similarly modulates hormone signalling in the body to promote cross-talk with other fats or other tissues.

Second, g-beige fat is induced through a previously unknown pathway, which might

be targeted specifically to improve glycolytic control — a key factor in the treatment of type 2 diabetes. And if these cells are also present in humans, their characterization might pave the way for new approaches to activating thermogenic fat cells. For such approaches to be successful, researchers would first need to determine how g-beige fat is induced, including which receptors and subsequent signalling cascades trigger GABPa-mediated differentiation.

In light of the current obesity epidemic, efficient and safe approaches are required to regulate excess body weight. The discovery of g-beige fat could provide a sweet way to do just that. ■

Wenfei Sun and Christian Wolfrum are at the Institute of Food, Nutrition and Health, ETH Zurich, 8603 Schwerzenbach, Switzerland. e-mail: christian-wolfrum@ethz.ch

1. Cannon, B. & Nedergaard, J. *Physiol. Rev.* **84**, 277–359 (2004).
2. Enerbäck, S. *Cell Metab.* **11**, 248–252 (2010).
3. Chen, Y. *et al. Nature* **565**, 180–185 (2019).
4. Rosen, E. D. & Spiegelman, B. M. *Cell* **156**, 20–44 (2014).
5. Cypess, A. M. *et al. Cell Metab.* **21**, 33–38 (2015).
6. Ye, L. *et al. Proc. Natl Acad. Sci. USA* **110**, 12480–12485 (2013).
7. Razzoli, M. *et al. Mol. Metab.* **5**, 19–33 (2016).
8. Schwallie, P. C. *et al. Nature* **559**, 103–108 (2018).
9. Burl, R. B. *et al. Cell Metab.* **28**, 300–309 (2018).
10. Villarroya, F., Cerejio, R., Villarroya, J. & Giralt, M. *Nature Rev. Endocrinol.* **13**, 26–35 (2017).

This article was published online on 19 December 2018.

EARTH SCIENCE

Experiments cast light on mantle mineralogy

A technically challenging analysis has revealed the physical properties of a mineral at pressures and temperatures as high as those in Earth's mantle. The findings have implications for our understanding of Earth's deep interior. SEE LETTER P.218

JOHANNES BUCHEN

The movement of tectonic plates carries parts of Earth's crust back into the deep interior of the planet. Earlier in Earth's history, processes that preceded the present regime of tectonic motion might also have forced crustal rocks to sink into Earth's mantle¹. The fate of crustal material recycled in the mantle, however, is unknown. On page 218, Gréaux *et al.*² present measurements of the speed of sound in one of the minerals that are thought to form in recycled crustal rocks, at pressures and temperatures that resemble those in Earth's mantle. The findings promise

to facilitate the use of seismic waves to track crustal rocks in the mantle.

Beneath the oceans, Earth's crust consists mainly of a rock called basalt. When exposed to high pressures as a result of tectonic processes, basalt transforms into an assemblage of minerals that is denser than the rocks of the underlying mantle. Basaltic rocks can therefore sink into the mantle, where they are exposed to rising pressures and temperatures that drive further mineral transformations^{3,4}. One of these reactions produces the mineral calcium silicate perovskite (CaPv).

The crystal structure of CaPv has cubic symmetry at temperatures greater than about