

process. Meanwhile, tot no. 2 scrambles up the east side. Dad plants tot no. 1 on the ground, and heads up the east side to retrieve tot no. 2, while tot no. 1 breezes back up the west side. When dad and tot no. 2 return to the bottom, one full cycle is complete. Of course, no self-respecting parent would put up with this nonsense, but what jet streams lack in intelligence, they make up in perseverance.

In fact, the troposphere is a playground for a large array of westward- and eastward-propagating waves that are constantly clambering up into the stratosphere. Forty years ago, this was recognized⁵ as being key to the QBO mechanism. Each type of wave has a different personality. Buoyancy waves (internal gravity waves) travel in all directions; they are just like the waves on the ocean, but internal to the atmosphere. Near the equator, the vorticity waves (Rossby waves) travel slowly westward. In addition, there is a class of 'half waves' (equatorial Kelvin waves) that lean against each other across the equator; they travel rapidly eastward.

When any of these waves drifts upwards and encounters a stratospheric jet going in the same direction, it deposits its momentum just shy of the jet maximum, which has the effect of coaxing the jet downwards in a slow but continuous manner; this led to the first complete description⁶ of the QBO mechanism. Meanwhile, waves that travel in the opposite direction are not blocked but can scramble all the way to the top of the climbing frame, thereby starting a new jet in their direction, which then slowly descends. The upshot is that the roughly 2-year period of the QBO on Earth is governed more by the strength of the wave flux, and the size and shape of the stratosphere, than by the rate of rotation or revolution of the planet.

Significantly, Jupiter also exhibits a QBO-like oscillation, with a period of 4.5 (Earth) years, appropriately called the QJO (quasi-quadrennial oscillation)⁷. As described by Fouchet *et al.*², similarities between Saturn's 15-year oscillations and the QBO/QJO include strong equatorial confinement of temperature extremes; asymmetry between the eastward and westward wind shears, with stronger eastward shears; and out-of-phase temperature changes between the equator and 15–20° latitude.

My guess is that the lengthening of the period from Earth to Jupiter to Saturn relates to their decreasing proximity to the Sun, which reduces the total energy budget available to their waves. But watch out: in addition to the QBO, Earth's stratosphere exhibits a strong signal with a 0.5-year period; this is the semiannual oscillation (SAO) mentioned earlier (Mars also seems to show an SAO⁸). The Saturn year is 29.5 Earth years, meaning that the Saturn wave is, at least descriptively, an SAO (for comparison, the corresponding time ratio for Jupiter's QJO is 4.5/11.9 = 0.38, which is not semiannual). Whereas Earth's SAO is driven by the different response to surface heating between the ice of Antarctica and the surrounding ocean,

it is not obvious what the analogy is on Saturn (although its large ring shadow probably has a role).

The influence of the QBO and SAO on Earth's weather cannot be overstated. They modulate seasonal activity, the behaviour of the Hadley cell (the overturning circulation that predominates in the tropics), the strength of the polar vortex, the mixing of atmospheric trace species, and even the predictability of regional patterns such as the Indian monsoon in August and September⁹. Because the portfolio of eastward waves is distinct from that of westward waves, the eastward and westward phases of the QBO are different. The big news is that this asymmetrical, long-period response has now been observed in the stratospheres of three planets. The question for modellers is whether these stratospheres are like three

different string instruments, or are more like three of the same instruments being played differently. ■

Timothy E. Dowling is at the Comparative Planetology Laboratory, University of Louisville, Louisville, Kentucky 40292, USA.
e-mail: dowling@louisville.edu

1. Orton, G. S. *et al. Nature* **453**, 196–199 (2008).
2. Fouchet, T. *et al. Nature* **453**, 200–202 (2008).
3. Baldwin, M. P. *et al. Rev. Geophys.* **39**, 179–229 (2001).
4. Andrews, D. G., Holton, J. R. & Leovy, C. B. *Middle Atmosphere Dynamics* (Academic, San Diego, 1987).
5. Lindzen, R. S. & Holton, J. R. *J. Atmos. Sci.* **25**, 1095–1107 (1968).
6. Plumb, R. A. *J. Atmos. Sci.* **34**, 1847–1858 (1977).
7. Leovy, C. B., Friedson, A. J. & Orton, G. S. *Nature* **354**, 380–382 (1991).
8. Kuroda, T., Medvedev, A. S., Hartogh, P. & Takahashi, M. *EPSC Abstr.* **2**, EPSC2007-A-00204 (2007).
9. Claud, C. & Terray, P. *J. Clim.* **20**, 773–787 (2007).

COMPUTATIONAL BIOCHEMISTRY

Old enzymes, new tricks

Giovanna Ghirlanda

Although enzymes are superb catalysts, their range of reactions is limited to those that support life. Their repertoire could be expanded by a method that allows artificial enzymes to be made from scratch.

Enzymes are astoundingly good catalysts: they allow reactions to occur billions of times faster than would be possible without them, at temperatures much lower than those required by typical synthetic catalysts. But enzymes have evolved to accelerate only biological reactions, under the narrow set of conditions that are compatible with life. Two papers from the same group, one in this issue (Röthlisberger *et al.*¹, page 190) and another in *Science* (Jiang *et al.*²), show how these limitations can be overcome. They describe a method for designing enzymes that catalyse unnatural reactions, and demonstrate its use for two different chemical transformations.

Enzymes work by lowering the activation energy of reactions, specifically by confining substrates in binding sites that stabilize the highest-energy arrangement of atoms in the reaction pathway (known as the transition state). They also shield the reactants, thus preventing possible side reactions. The idea behind the latest work^{1,2} is simple — model the transition state for a reaction, stabilize it by surrounding it with carefully placed chemical groups, graft the resulting active site into an existing protein and then alter the amino-acid sequence of the protein to accommodate the changes. In practice, this is a complicated procedure. For starters, building an accurate model of a transition state requires a detailed understanding of the reaction's mechanism, which isn't always available. Furthermore, transition states are modelled using quantum-mechanical calculations, but

currently available methods can handle only a limited number of atoms, and are often inadequate for modelling enzyme reactions.

Designing a protein that folds into a given structure is equally challenging. For a protein made of 100 amino acids, there are about 10¹³⁰ possible sequences, each of which can adopt many different conformations. The thermodynamic stability of every sequence and conformation must therefore be calculated to find the lowest-energy structure (that is, the one most likely to form). Some simplifications can be made using advanced computational methods to quickly eliminate unfavourable combinations. This has resulted in several notable accomplishments, such as the complete redesign of a protein consisting of 28 amino acids³, the design of an amino-acid sequence that forms a structure not found in nature⁴, and the engineering of naturally occurring proteins into biosensors for trinitrotoluene (TNT) and other small molecules⁵.

With these precedents, you might think that designing catalytic proteins should be straightforward, but success has been limited. Catalytically inactive proteins have been converted into modestly catalytic ones for two different reactions, but the observed enhancements of rate^{6,7} were only about a millionth of those produced by naturally occurring enzymes. It is also sometimes difficult to prove that designer enzymes are truly catalytic on the basis of biochemical observations, and some exciting claims have been found to be flawed.

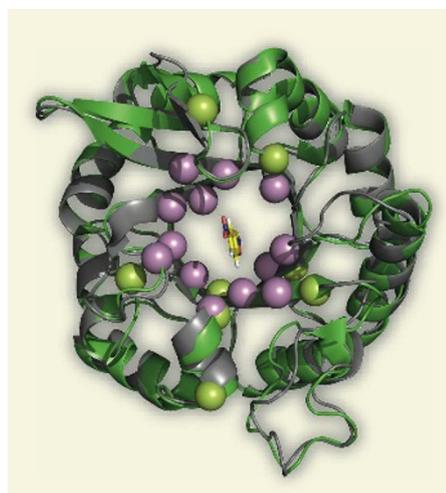


Figure 1 | Enzymes by design. Röthlisberger *et al.*¹ have computationally designed and prepared the first enzyme capable of catalysing a non-biological reaction. Here, the computational model (grey) is overlaid with the crystal structure of the actual protein (green); the two overlap almost perfectly. The substrate is shown at the centre of the structure. The design process involved modifying the amino-acid sequence of a naturally occurring protein. Residues selected computationally to form the active site are shown as purple spheres. Additional mutations that were introduced *in vitro* to optimize the enzyme's performance are shown as green spheres.

But some reports of catalysis by designed enzymes have fared rather better — especially those that are based on sound crystallographic evidence^{6–8}. An essential step in demonstrating the success of a designer enzyme, therefore, is the determination of a high-resolution crystal structure for the protein, to verify that the designed catalytic features are present. The results of Röthlisberger *et al.*¹ and Jiang *et al.*² are remarkable in the spectacular agreement between their computationally predicted enzyme models and the experimentally determined structures (Fig. 1).

Röthlisberger *et al.*¹ made an enzyme that catalyses the Kemp elimination reaction (see Fig. 1a on page 190 for a reaction scheme). The Kemp elimination is initiated by the removal of a hydrogen ion from a carbon–hydrogen bond in the substrate; the minimum requirement for catalysis of the reaction is the presence of a base to perform this step. The authors therefore identified two amino acids — aspartic acid and histidine — that have side chains that can act as bases under physiological conditions, and used these as the starting points of their putative active sites. They decorated models of the proposed active sites with other chemical groups found in proteins, choosing those that could interact favourably with groups in the substrate. They then used state-of-the-art quantum-mechanical methods to precisely place all the groups in the models to maximize stabilization of the transition state of the substrate. The authors thus obtained a large ensemble of designs for catalytic sites in enzymes.

Next, Röthlisberger *et al.* selected about 100 proteins that could be used as scaffolds for their proposed active sites. The criteria for selection were the availability of high-resolution crystal structures and the presence of pre-organized cavities, with a preference for proteins that behave well in experiments (that is, those that have good solubility, are expressed easily in cells, and so on). The authors then used computational methods to search each of the proteins for specific regions that could accommodate the sites, narrowing down the vast number of possibilities to about 100,000 promising leads. These were whittled down further using an automated modelling technique to find the optimal amino-acid sequence in defined shells around the active site, selecting sequences that maintained protein stability and integrity.

This computational screening method picked out 59 candidate enzymes, which the authors expressed in cells and evaluated for their ability to catalyse the target reaction. Only eight of the proteins had measurable catalytic activity. The team then used *in vitro* evolution to further optimize one of their successful leads (designated KE07), mutating the amino-acid sequence in both random and directed locations. After several rounds of mutation and screening, Röthlisberger *et al.* obtained improved enzymes that were up to 200 times more active than KE07. The best two of these mutants accelerate the rate of the Kemp elimination reaction to about a million times that of the uncatalysed version.

The strategy used by Röthlisberger *et al.*¹ promises to be general, as the same group² has successfully applied the procedure to another chemical transformation known as the retro-aldol reaction, which is very different from the Kemp elimination. The complexity of the design procedure is underlined by the number of interdisciplinary groups involved in the work, and by the huge amount of computational power required to solve the problem — donated from hundreds of thousands of idling computers around the world as part of a project known as Rosetta@home⁹.

Those in the know might say that the performance of the designed enzymes is far from impressive — the reaction-rate enhancements for typical, naturally occurring enzymes are anywhere between 10,000 and 1 billion times higher than those of the artificial enzymes described in these papers^{1,2}. Furthermore, the chosen reactions are relatively easy targets. The Kemp elimination is accelerated by several catalysts, including various synthetic compounds, catalytic antibodies and even serum albumin. Similarly, the retro-aldol reaction is catalysed by antibodies¹⁰ and by various peptides^{11,12}. Indeed, the rate enhancements reported by Röthlisberger *et al.*¹ are equivalent to those of only the most sophisticated catalytic antibodies^{13,14}; the enhancements obtained by Jiang *et al.*² for the retro-aldol reaction are even more modest.

Another limitation of the design process is that, although naturally occurring enzymes



50 YEARS AGO

A Hundred Years of

Evolution. By Dr. G. S. Carter

— It is fundamental to the neo-Darwinian theory that Weismann's concept of the inviolability of germ plasm by soma is correct, and that mutational changes in the gene complex arise solely at random; Dr. Carter (p. 87) accepts Weismann's doctrine as "undeniable when once pointed out". It is arguable, however, that the "separateness of the gonad from the rest of the soma" is a philosophical concept of the same order as that of the soul and the body. As such it may have been valid in the state of biological knowledge in Weismann's time, but it has to-day become undermined to the point of collapse ... That mutation is random is purely theoretical, depending in the first place on the validity of the divorce between germ plasm and soma, and in the second upon the absence of evidence to the contrary.

From *Nature* 10 May 1958.

100 YEARS AGO

To be told that life exists on Mars tells us but little of its nature

... Perhaps on Mars there is only one living being, a gigantic vegetable the branches or pseudopodia of which embrace the planet like the arms of an octopus, suck water from the melting polar snows, carry it to other parts of the planet, and are visible to us as the Martian canals. Lowell adduces the straightness of the canals as a proof that they are artificial products of intelligent beings. But they are certainly no straighter than the somewhat similarly interlaced pseudopodia seen in certain Heliozoa, Foraminifera, and Radiolaria ... My position is that one may admit that Prof. Lowell's brilliant researches prove the existence of life on Mars, and still ask from him further evidence before we are convinced that that life is intelligent.

From *Nature* 7 May 1908.

50 & 100 YEARS AGO

have evolved to optimize steps other than just catalysis (such as the binding of substrates and the release of products), the model used by the authors^{1,2} to design their enzymes doesn't attempt to address these factors. This is understandable, because many of the finer features that provide enzymes with their unique properties are not yet understood. For example, the mutations introduced by the authors into their enzymes by directed evolution did not modify the active site itself, but occurred at neighbouring positions (Fig. 1). The effect of some of these mutations can be easily understood with hindsight, but others are much less obvious. It was therefore wise of the authors to let nature lend a helping hand in their designs.

Nevertheless, these results^{1,2} are a milestone

in biochemistry. For the first time, artificial enzymes have been designed for non-biological reactions, providing rate accelerations that are about 1,000 times faster than previous examples of computationally designed enzymes. Biochemists have long wanted to build artificial enzymes to identify and validate the minimal requirements for enzyme-like catalysis. These reports provide an accurate framework for this enterprise to which further features can be added. As Röthlisberger *et al.*¹ note, the ability to design enzymes will truly test our understanding of enzyme catalysis. ■

Giovanna Ghirlanda is in the Department of Chemistry and Biochemistry, Arizona State University, Tempe, Arizona 85287-1604, USA. e-mail: giovanna.ghirlanda@asu.edu

1. Röthlisberger, D. *et al.* *Nature* **453**, 190–195 (2008).
2. Jiang, L. *et al.* *Science* **319**, 1387–1391 (2008).
3. Dahiyat, B. I. & Mayo, S. L. *Science* **278**, 82–87 (1997).
4. Kuhlman, B. *et al.* *Science* **302**, 1364–1368 (2003).
5. Looger, L. L., Dwyer, M. A., Smith, J. J. & Hellinga, H. W. *Nature* **423**, 185–190 (2003).
6. Bolon, D. N. & Mayo, S. L. *Proc. Natl Acad. Sci. USA* **98**, 14274–14279 (2001).
7. Kaplan, J. & DeGrado, W. F. *Proc. Natl Acad. Sci. USA* **101**, 11566–11570 (2004).
8. Geremia, S. *et al.* *J. Am. Chem. Soc.* **127**, 17266–17276 (2005).
9. <http://boinc.bakerlab.org/rosetta/>
10. Sinha, S. C., Barbas, C. F. & Lerner, R. A. *Proc. Natl Acad. Sci. USA* **95**, 14603–14608 (1998).
11. Tanaka, F. & Barbas, C. F. *Chem. Commun.* 769–770 (2001).
12. Tanaka, F., Fuller, R. & Barbas, C. F. *Biochemistry* **44**, 7583–7592 (2005).
13. Hu, Y., Houk, K. N., Kikuchi, K., Hotta, K. & Hilvert, D. *J. Am. Chem. Soc.* **126**, 8197–8205 (2004).
14. Thorn, S. N., Daniels, R. G., Auditor, M.-T. M. & Hilvert, D. *Nature* **373**, 228–230 (1995).

DEVICE PHYSICS

Chance match

Robert M. Westervelt

A clever device uses the quantum statistics of electron tunnelling to match image patterns. The circuit is low-power, works at room temperature — and could point to a way forward for silicon electronics.

Over the past three decades, as the components that make up integrated circuits have been made smaller and smaller, the power of computer chips has grown exponentially, even as their cost has fallen drastically. But sooner rather than later — by around 2020, according to one estimate¹ — the scaling-down process will become difficult to maintain^{2,3}.

The energy required to represent a bit of information will become larger than the heat that can be carried away from a tiny circuit element; what's more, as devices approach the size of atoms, quantum-physical phenomena will become important, changing even the ground rules of how bits are processed. Writing in *Applied Physics Letters*⁴, Nishiguchi *et al.* detail what might be one way to circumvent, and even exploit, these issues. They describe a circuit that allows them to perform the computing operation of pattern matching by harnessing the stochastic, quantum-mechanical tunnelling of single electrons into a transistor⁵.

Pattern recognition is a natural enough task for people, but is often difficult for computers. We would like our computer processors to be like us and recognize an object (a cat or a dog, say) in a photographic image, understand the meaning of spoken words, or translate efficiently from one language to another. But pattern recognition also has more abstract, fundamental

uses: extracting a simple conclusion from a great body of input data, for instance.

Nishiguchi *et al.*⁴ build on previous work^{6–8} to construct a simple pattern-matching circuit using a basic building-block of two transistors (more precisely, metal–oxide–semiconductor field-effect transistors, or MOSFETs) pat-

terned on a silicon-on-insulator wafer. They trap and store single electrons on the first of these nanoscale transistors, the 'T-FET'. They are able to reduce the rate at which electrons tunnel quantum-mechanically into a storage node on the T-FET to very low levels of around one per second. The authors show that the trapped electrons obey Poisson statistics, and represent a statistically random source of events that can be used for stochastic signal processing⁹.

The job of the second transistor that makes up the authors' processor, the 'D-FET', is to detect the number of electrons stored in the T-FET. It does this through a capacitive coupling: as the number of electrons stored in the T-FET increases, the current passing through the D-FET is progressively reduced. The coupling is sensitive enough that the tunnelling of a single electron into the T-FET is registered as a discrete drop in current in the D-FET.

To perform pattern matching, the individual bits of an input image must be compared with those of a reference image. Nishiguchi and colleagues set an input bit, i , to 0 or 1 by stepping the input 'source' voltage of the T-FET. Similarly, they set a reference bit, r , by changing the T-FET's 'gate' voltage, which controls the passage of current from the source into the storage node (Fig. 1). When $i = 0$, the tunnelling rate into the T-FET is negligible. When $i = 1$, tunnelling occurs, and the number of electrons stored in the T-FET slowly builds up. The precise rate of this tunnelling is controlled by the gate voltage, and thus the reference bit: it is large when $r = 1$, and small when $r = 0$.

Essentially, this set-up creates a detector that flags up when the input and reference bits are both on, $i = r = 1$: in this case, electrons build up in the T-FET particularly

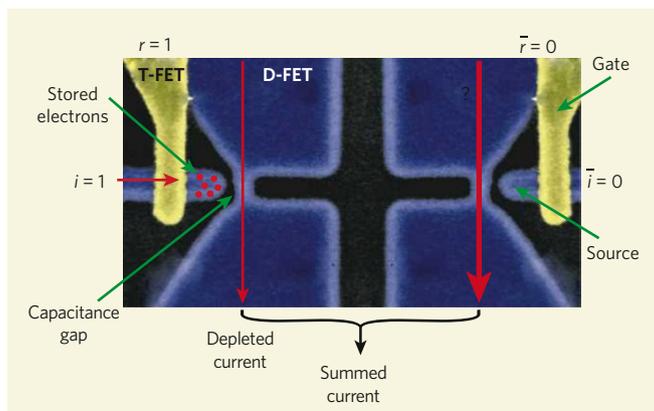


Figure 1 | Dual processor. Nishiguchi and colleagues' pattern-recognition processor⁴ uses two basic components that each consist of two capacitatively coupled transistors: a transfer transistor (T-FET) and a detector transistor (D-FET). The probability that an electron will tunnel from the source of the T-FET, under the gate and into the storage node is determined by the source voltage, which is set by the value of a bit i in the input image, and by the gate voltage, which is set by a bit r in the reference image. The more electrons accumulate in the T-FET storage node, the lower the current that flows through the capacitatively coupled D-FET. In the instance depicted, both the input and reference bits are turned on, $i = r = 1$, and electrons accumulate in the storage node, reducing the detector current. The second unit (right) is fed with the inverse inputs of the first, \bar{i} and \bar{r} . If the original inputs were matched at 0, the inputs here would be 1, and this half of the processor would record the depleted current characteristic of matched bits. (Figure adapted from ref. 1.)