peratures and is partially trapped by the magnesia, with which it is intertwined, and so it does not move significantly during this reaction. Consequently, the original shape of the silica is largely preserved. The magnesia can be selectively removed by bathing the solid in hydrochloric acid, leaving behind a silicon replica of the original silica structure.

Bao and colleagues demonstrate their approach with several types of diatom. Because of their abundance, these unicellular, photosynthesizing algae are important in soil, freshwater and seawater ecosystems. Diatoms fortify their cell walls with silica, and so have also been used to study biomineralization[7]. Their exoskeletons exhibit a variety of beautifully intricate, species-specific shapes and patterns (Fig. 1).

The authors use their technique to convert these silica skeletons into silicon. Their results show that the overall shapes, with their intricate pores and channels, are conserved. Also, as oxygen is removed from the structure, fewer atoms remain. This introduces new, nanometre-scale pores into the replica. For some applications — those that require solid silicon — this would be a disadvantage. For others it can be desirable for two reasons. First, the presence of the pores significantly increases the surface area of the final material, which can be important for applications such as sensing. Second, the final structure is composed of nanoscale crystals of silicon, which have useful optical properties. In particular, unlike bulk silicon, nanocrystalline silicon can fluoresce efficiently[8,9].

Bao et al.[5] test both of these effects. First, they attach wires to a single diatom replica and use it as a microsensor for nitric oxide. Their results suggest that the diatom-derived silicon structure can provide a much more efficient sensor element than other, more conventional approaches. Second, they find that their silicon replicas can fluoresce after being partially oxidized in water.

This work indicates that many different silicon materials can be derived from diatoms, as well as other silica structures harvested from nature. More generally, it complements other recently developed 'templating' techniques. Solids structured on nanometre or micrometre length scales can now be obtained through a variety of simple self-assembly routes. These templates are then filled with a different material of interest before the template is removed. Silica templates are common because they are so easy to prepare, and silicon has been used as an infill material[10] because of its technological importance and its high refractive index, which is useful in some applications in photonics.

A templating approach obviously produces the 'negative' of the original template. To obtain a duplicate of the template, the procedure must be repeated[11]: the negative structure is used as a mould, infiltrated a second time and removed. The mould material must be chosen carefully for all of these steps to be completed successfully. In contrast, Bao et al.[5] have provided a

method for copying a structure into silicon directly, albeit with the introduction of nanometre-scale pores. In combination with other self-assembly tricks, this opens the way to a variety of exciting new fabrication strategies. The implications might encompass not just biologically derived materials, but also many man-made creations. ■

David J. Norris is in the Department of Chemical Engineering and Materials Science, University of Minnesota, 421 Washington Avenue SE, Minneapolis, Minnesota 55455, USA. He is currently on sabbatical at the Walter Schottky Institute, Garching, Germany.
e-mail: dnorris@umn.edu

1. Vukusic, P. & Sambles, J. R. Nature 424, 852–855 (2003).
2. Aizenberg, J. et al. Science 309, 275–278 (2005).
3. Heuer, A. H. et al. Science 255, 1098–1105 (1992).
4. Mann, S. & Ozin, G. A. Nature 382, 313–318 (1996).
5. Bao, Z. et al. Nature 446, 172–175 (2007).
6. Iler, R. K. The Chemistry of Silica (Wiley, New York, 1979).
7. Sumper, M. Science 295, 2430–2433 (2002).
8. Canham, L. T. Appl. Phys. Lett. 57, 1046–1048 (1990).
9. Littau, K. A., Szajowski, P. J., Muller, A. J., Kortan, A. R. & Brus, L. E. J. Phys. Chem. 97, 1224–1230 (1993).
10. Chomski, E. & Ozin, G. A. Adv. Mater. 12, 1071–1078 (2000).
11. Jiang, P., Bertone, J. F. & Colvin, V. L. Science 291, 453–457 (2001).

EVOLUTIONARY BIOLOGY

The Elvis paradox

Andrew Hendry

Evidence for a universal driver of evolution across all timescales could mean that the venerable paradox of stasis is dead. But even with such evidence, some biologists would be reluctant to accept its passing.

Disagreement has long swirled around the relative importance of various forces that might drive evolution on timescales ranging from dozens to millions of generations. Writing in *The American Naturalist*, Estes and Arnold[1] offer a provocative contribution to this debate: they propose that evolutionary changes on *all* timescales might be explained by a single, simple model of adaptation.

Much of the challenge can be distilled down to what has been called the 'paradox of stasis'[2]. For me, the most obvious manifestation of this paradox is that neo-darwinian theory, with its emphasis on the power of selection, predicts the potential for rapid adaptation, whereas most lineages of organisms instead show long-term stasis: that is, very little cumulative change over long periods of time[3,4]. Several hypotheses have been advanced in the hope of resolving this seeming discontinuity between short- and long-term evolution[2,4,5], but none has been convincing enough to resonate across the various camps.

Estes and Arnold[1] point out that the best way to discriminate between the hypotheses is to confront the predictions of alternative evolutionary models with the reality of data. This sort of comparison has recently been made possible by compilations of data on phenotypic changes (such as in mean body size) within animal lineages at a variety of different timescales[4,6,7]. The pattern emerging from these data is that phenotypic changes over dozens of generations can range from small to large, and that this range remains roughly the same even over millions of generations (Fig. 1). This pattern thus affirms the original paradox — that phenotypic change can be dramatic on short timescales, but rarely accumulates into substantial evolutionary trends.
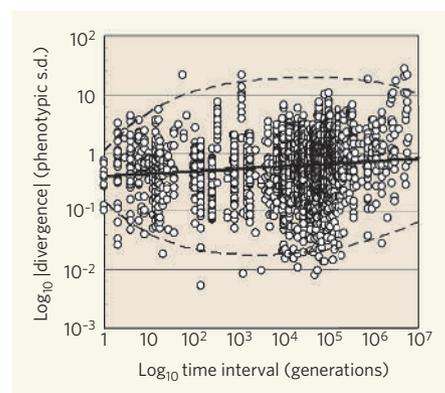
Figure 1 | The database of phenotypic changes analysed by Estes and Arnold[1]. The graphic summarizes phenotypic changes (usually body size or other aspects of morphology) through time within animal lineages, as well as phenotypic differences between lineages that had a common ancestor at a known time in the past. The lineages concerned include a wide variety of vertebrate and invertebrate taxa, ranging from snails to flies to salmon to mice to horses. Dotted lines are the 99% confidence ellipse for the data. The solid line is the regression line through the data. (Reproduced from ref. 1.)

Estes and Arnold[1] evaluate the degree to which six evolutionary models fit the observed data. All of these models are based on 'adaptive landscapes' — an analytical framework that relates mean phenotypes (mean body size, for instance) to the expected mean fitness (that is, number of offspring) of a population (Box 1, overleaf). Evolution on such landscapes tends towards 'hill climbing', where the mean phenotype of the population moves towards that which maximizes population fitness (a local fitness peak)[8,9].

Three of the models tested by Estes and

Box 1 | Evolution on six adaptive landscapes

Adaptive landscapes, shown here in blue, depict the mean fitness of individuals in a population (for example, average number of offspring) for a range of possible mean phenotypes (for example, average body size) in that population. A population with a particular phenotypic distribution (an example is shown in red) will tend to evolve up the steepest adjacent slope on the adaptive landscape until its mean approaches that which maximizes population fitness. Estes and Arnold[1] consider evolutionary changes that might result from several different types of adaptive landscape.

Under 'neutrality', the adaptive landscape is flat, without any fitness peaks, and so the mean phenotype will drift around at random (red arrows). Under

'brownian motion', the adaptive landscape has a single peak, the location of which drifts around at random (blue arrows). Under this model, the change in peak position is calculated from the position that the peak reached in the previous generation. Under 'white-noise motion', the peak again moves at random, but this time the change in each generation is calculated from an unchanging starting point (that is, the changes are not additive). For these models of random peak movement, mean phenotypes will follow the randomly moving peak.

Under 'moving optimum', the adaptive landscape has a single peak, and this peak is moving predictably in a particular direction every generation, with mean phenotypes following after. Under 'peak shift', the adaptive landscape has two fitness peaks separated
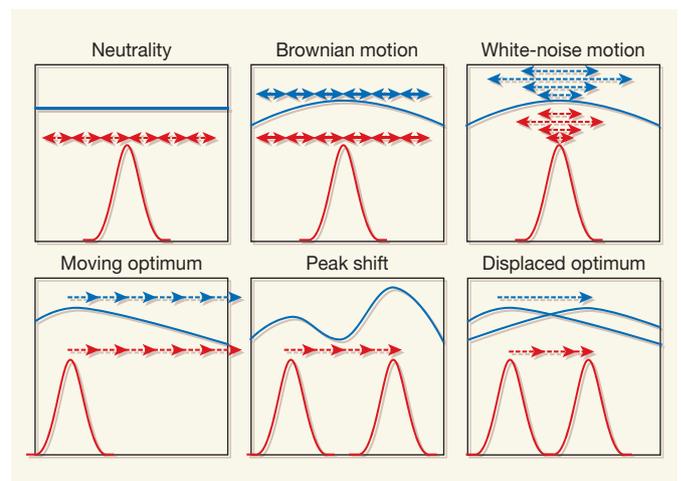
by a valley, and mean phenotypes may sometimes shift from the lower peak to the higher one. Under 'displaced optimum', the adaptive landscape has a single peak, and this peak jumps in a

single generation to a new location, and then stays there. The mean phenotype then evolves over to the new peak. It is this model that Estes and Arnold find best fits the data. **A.H.**

Arnold represent different flavours of randomness (Box 1). One model ('neutrality') evokes a flat adaptive landscape, on which the mean phenotype of a finite population will drift at random. The other two random models have a single fitness peak that moves randomly according to either 'brownian motion' or 'white-noise motion'. In these two models, mean phenotypes forever chase the randomly moving fitness peaks, like a new task for Sisyphus. Estes and Arnold argue that these three models fail to fit the data well, suggesting that randomness, at least in these forms, may not be a primary driver of phenotypic change. This conclusion will be reassuring, or perhaps just obvious, to the innumerable evolutionary biologists who believe that adaptation plays a central role in evolution.

The other three models involve a directional shift in the position of a fitness peak. In one ('moving optimum'), the adaptive landscape has a single peak that moves step-by-step in a particular direction, with the phenotypic mean of the population following along. This model predicts too little evolution on short timescales and too much on long timescales relative to the observed data. In another directional model ('peak shift'), the adaptive landscape has two peaks and, under some conditions, the phenotypic mean can shift rapidly from one peak to the other. This model predicts too little evolution on short timescales, and it only fits the data well on long timescales when populations are unrealistically small. In the final model ('displaced optimum'), the adaptive landscape has only one peak, and the position of this peak shifts abruptly, but just once. In the authors' estimation, this last model fits the data quite well.

The displaced-optimum model can be visualized by reference to a hypothetical population

that is well adapted to its local environment, in which case the mean phenotype of the population (let's say a body mass of 10 g) will be centred near a local fitness peak (also 10 g). Then imagine that the environment changes abruptly and displaces the optimum to a new location of 12 g, leaving the population mean behind at 10 g. Now the largest individuals in the population will be favoured by natural selection, and the mean phenotype will increase across generations until it is positioned at the new fitness peak (that is, 12 g), where it will then stay in the absence of further environmental change. This all makes sense, but the surprising part is that the displaced-optimum model assumes that this happens only once for a given lineage — regardless of timescale. The key general point, however, is that the peak can be displaced only a restricted amount (that is, within defined bounds), even if it takes several steps to get there.

Conveniently, Estes and Arnold[1] provide an Excel file in which every parameter in every model can be varied and the resulting outcomes compared with actual data. After playing with these models myself, I tend towards general agreement with the authors, adding the caveat that achieving the observed changes on short timescales requires a very large displacement of the optimum, coupled with a very sharp fitness peak. These properties mean that a population will have a substantial fitness decline immediately after the optimum moves — a possible recipe for extinction. Perhaps the largest changes on short timescales are the result of phenotypic plasticity (when a genotype expresses a different phenotype in a new environment), rather than of genetic change, the latter being the focus of the models. Other factors that may inflate short-term changes are sampling errors, conflation of

geographical variation with lineage evolution, and publication biases (perhaps published studies on short timescales tend towards those that find the largest changes).

Have Estes and Arnold[1] slain the paradox of stasis with a simple displaced-optimum model? In my opinion, the paradox might have been slain only in a broad sense, because the pattern in the data also might be replicated by other models that generate a range of short-term changes that do not accumulate into long-term trends. Such models might include various types of fluctuating selection, where fitness peaks on adaptive landscapes move back and forth owing to environmental variation[10]. Perhaps the paradox of stasis will have its final death at the point of a rapier, whereas Estes and Arnold have wielded a scimitar.

Whatever the model, it will have to generate rapid changes on short timescales, and yet still be constrained by boundaries on long timescales[4]. It is also possible that the paradox is a phantom, against which swords are of no use. Indeed, it may have been dead on arrival: way back in 1944, George Gaylord Simpson[8] suggested that evolutionary stasis might be explained by 'adaptive zones', where fitness peaks move back and forth within constrained bounds.

For some, any report of the death of this paradox will probably evoke the same reaction as the death of Elvis, with a large number of fans reluctant to accept its passing. But in the end, evolutionary biologists will probably converge on more pertinent questions, such as 'What generates and maintains adaptive zones in the first place?', and 'How do some lineages ultimately bridge the gap between different adaptive zones?'. This convergence would probably both please and frustrate Simpson, were he still alive, given that he posed much

the same questions more than 60 years ago. ■
Andrew Hendry is in the Redpath Museum and Department of Biology, McGill University, Montreal, Quebec H3A 2K6, Canada.
e-mail: andrew.hendry@mcgill.ca

1. Estes, S. & Arnold, S. J. *Am. Nat.* **169,** 227–244 (2007).
2. Hansen, T. F. & Houle, D. in *Phenotypic Integration* (eds Pigliucci, M. & Preston, K.) 130–150 (Oxford Univ. Press, 2004).
3. Gould, S. J. & Eldredge, N. *Paleobiology* **3,** 115–151 (1977).
4. Gingerich, P. D. *Genetica* **112–113,** 127–144 (2001).
5. Charlesworth, B. *et al. Evolution* **36,** 474–498 (1982).
6. Hendry, A. P. & Kinnison, M. T. *Evolution* **53,** 1637–1653 (1999).
7. Kinnison, M. T. & Hendry, A. P. *Genetica* **112–113,** 145–164 (2001).
8. Simpson, G. G. *The Tempo and Mode of Evolution* (Columbia Univ. Press, 1944).
9. Arnold, S. J. *et al. Genetica* **112–113,** 9–32 (2001).
10. Grant, P. R. & Grant, B. R. *Science* **296,** 707–711 (2002).

PLANETARY SCIENCE

Water cycling on Mars

Victor R. Baker

The Meridiani Planum region on Mars is rich in minerals derived from evaporation, but lacks a topography consistent with standing water. Do the deposits stem from upwelling groundwater early in the planet's history?

A succession of sophisticated spacecraft missions has led to spectacular advances in the understanding of Mars' global hydrology over the past few decades. One of many examples is the discovery of abundant hydrated sulphate salt minerals. These minerals are found at many locations on the planet — most notably at Meridiani Planum, the landing site of NASA's robotic Mars rover Opportunity — and prove that water must once have been abundant on the surface of Mars. On page 163 of this issue, Andrews-Hanna *et al.*[1] use a numerical model to simulate the evolving global flow of sub-surface groundwater early in Mars' geological history. They place their simulation in the context of the formation of the enormous

volcanic uplift feature known as Tharsis.

One way of developing a model of martian hydrology comes from a comparison with what we know about Earth. Western science was painfully slow in achieving its understanding of Earth's hydrological cycle. Many, if not most, of Isaac Newton's scientific contemporaries held the view that Earth's rivers were ultimately fed from upland springs. The springs were presumed to discharge water from within the planet, in much the same way as blood flows from a cut in an artery of the human body. Water from the oceans was presumed to return to the land through subsurface veins.

By contrast, Eastern philosophical writings had long held that Earth's water flowed as

Figure 1 | A well run dry? The Burns formation of Meridiani Planum is rich in evaporites. It is one of the sites where groundwater came to the surface and evaporated early in Mars' history, according to a model of the planet's hydrography produced by Andrews-Hanna *et al.*[1]. Earlier evidence from the Mars rover Opportunity for past evaporation of liquid water at the site had been difficult to explain in the absence of a topographic basin where water could stand. In the right foreground lies the Wopmay rock, whose distinctive lumpy appearance could be due to cracking processes caused by exposure to water.

NASA/JPL/CORNELL