with photons' or 'not imaging the hair at all'[4,5].

One might think that the same Zeno trick could be used to improve counterfactual quantum computers, so that they could actually be useful for something — or at least, that the probability of their finding the right answer should be better than that achieved by flipping coins or tossing dice. But this hope proved misguided: the straightforward application of the Zeno idea to counterfactual quantum computing did not seem to beat the random-guessing limit, and things were at a standstill.

Enter Hosten *et al.*[1]. These authors have discovered that chaining Zeno boosters together produces a super Zeno booster that is indeed capable of beating the random-guessing limit in counterfactual quantum computing. They also demonstrate that the idea works with a quantum-optical implementation of the Grover search algorithm. This algorithm proves that a quantum computer can always perform better than any classical computer in searching an unsorted database for a target datum[7]. It has a particularly simple implementation in quantum optical interferometry, as does the super-sized quantum Zeno effect. Hosten and colleagues show how the two can be married to locate the target element in a database — without running the quantum computer that searches for that element in the first place.

So what is counterfactual quantum computing good for, other than to illustrate the conclusion of a rather obscure quantum paradox? The computer must still be programmed and turned on, even if it is not run, so the approach will not save on electricity bills or labour costs. In fact, the value of the experiment lies simply in furthering our understanding of quantum mechanics and its interface with computation. The entire field of quantum information arose from physicists trying to understand the implications of paradoxes in quantum theory, and already the field is beginning to evolve on its own. The first commercial applications of quantum information, such as real-world quantum cryptography, are now being deployed.

A century ago, we saw the start of the first quantum revolution — the discovery of the quantum rules that underpin our world. We are now on the verge of a second revolution[8], in which these rules spawn technological applications. Results such as those of Hosten and colleagues[1] are significant markers on the road to that revolution. ■

Jonathan P. Dowling is at the Hearne Institute for Theoretical Physics, Department of Physics and Astronomy, Louisiana State University, Baton Rouge, Louisiana 70803, USA, and the Institute for Quantum Studies, Texas A&M University.
e-mail: jdowling@lsu.edu

1. Hosten, O., Rakher, M. T., Barreiro, J. T., Peters, N. A. & Kwiat, P. G. *Nature* **439,** 949–952 (2006).
2. Elitzur, A. C. & Vaidman, L. *Found. Phys.* **23,** 987–997 (1993).
3. Mitchison, G. & Jozsa, R. *Proc. R. Soc. Lond. A* **457,** 1175–1193 (2001).
4. Kwiat, P. *et al. Phys. Rev. Lett.* **74,** 4763–4766 (1995).
5. Kwiat, P., Weinfurter, H. & Zeilinger, A. *Sci. Am.* **275**(5), 72–78 (1996).
6. Knight, P. *Nature* **344,** 493–494 (1990).
7. Grover, L. K. *Phys. Rev. Lett.* **79,** 325–328 (1997).
8. Dowling, J. P. & Milburn, G. J. *Phil. Trans. R. Soc. Lond. A* **361,** 1655–1674 (2003).

NEUROBIOLOGY

Efficiency measures

Michael R. DeWeese and Anthony Zador

The nervous system translates sensory information into electrical impulses. The neural 'code' involved seems to represent natural sounds and images efficiently, using the smallest number of impulses.

Our perception of the outside world relies on the transformation of physical signals (such as light and sound) into a pattern of neural impulses, or spikes. These spikes are then transmitted to higher brain regions, where they are further transformed into other patterns of sensory spikes, and ultimately into the motor spikes that mediate behaviour. What is the relationship (the 'neural code') between these neural responses and the sensory signals they represent? Are there general principles underlying the neural code?

The 'efficient-coding hypothesis'[1] proposes that sensory neurons are adapted to the statistical properties of sensory signals to which neurons are exposed. Two papers in this issue invoke this principle to predict how neurons encode natural auditory and visual stimuli, as opposed to the artificial stimuli often used in experiments. Smith and Lewicki (page 978)[2] develop an algorithm to find an efficient representation of natural sounds and speech, and show that this theoretically predicted representation matches that observed experimentally in the auditory nerve of cats. Sharpee and colleagues (page 936)[3] show that cortical neurons adapt over seconds or minutes during the course of an experiment to maximize the information they provide about the stimulus. Together, the two papers show how the efficient-coding hypothesis can help to make sense of properties of the neural code on both evolutionary and behavioural timescales.

In the cochlea, sound is encoded into spikes, which are transmitted along the auditory nerve to higher stations in the auditory system. Auditory nerve fibres each respond to a narrow range of sound frequencies, with the range generally increasing with the median frequency. The response of each auditory nerve fibre can therefore be modelled as a (nonlinear) 'filter' that removes frequencies outside a particular range. Why do the auditory nerve filters have the particular form they do? Smith and Lewicki reasoned that if the auditory code is indeed 'efficient', then they should be able to predict the form of the auditory filter bank by finding the sparsest code; that is, the one that requires the least activity.

To obtain this prediction, Smith and Lewicki first expressed the efficient-coding hypothesis as an algorithm whose input is an ensemble of sounds, and whose output is a sparse encoding for transmitting or representing this ensemble. The algorithm discovers that the sparsest encoding of sounds is into brief events suggestive of spikes, the precise timing of which conveys much of the information. The sparsest code depends on the ensemble of sounds to be encoded; a code that is most efficient for one set of sounds is not necessarily most efficient for another.

Why should the most efficient code depend on the stimulus ensemble? The basic intuition is straightforward. Suppose I ask you to describe individual sounds produced by different musical instruments, but I limit your vocabulary to only four words (of your choosing). If you know that the instruments are used in a rock band (that is, they are chosen from the rock ensemble), you might choose a code consisting of the words 'guitar', 'bass', 'drums', 'keyboard'; but if the instruments are used in a classical orchestra (the classical ensemble), you might choose instead 'woodwind', 'brass', 'percussion', 'string'. So, the choice of the most efficient code depends on what is being described.

The dependence of the sparsest code on the ensemble raises the question of exactly which ensemble should be considered the evolutionarily relevant 'natural' ensemble. Smith and Lewicki therefore tested three subsets of sounds: animal vocalizations, transient environmental sounds such as crunching leaves and cracking twigs, and ambient environmental sounds like rain. They found that if the ensemble consisted of any single sub-ensemble, then the predicted sparse code failed to match previous experimental observations of cat auditory filters. However, the code predicted from a carefully balanced mixture of these natural sounds provided a good match to the experimental data, supporting the idea that these filters evolved to encode natural sounds efficiently. Interestingly, filters calculated from a human speech ensemble also fit the cat data well; because (as any cat owner will attest) the cat auditory system did not evolve to

©2006 Nature Publishing Group

attend to human speech, the authors speculate that speech may have evolved to match the properties of cochlear filters.

The efficient-coding hypothesis posits that the neural code should be adapted to the statistical properties of stimuli on all timescales. Where Smith and Lewicki were studying adaptation that had occurred over an evolutionary timescale, Sharpee and colleagues[3] looked for adaptation over seconds and minutes, the much faster timescale relevant to behaviour.

Sharpee and colleagues studied adaptation to natural scenes (images of a forest) of a particular subclass of neurons — the 'simple cells'— in the primary visual cortex of anaesthetized cats. Simple cells respond to oriented bars or edges, and it was already known that they could adapt to the most basic statistical properties of images — the luminance, mean and contrast (variance). So the authors probed these neurons with a control ensemble consisting of white-noise images (like the static on an untuned television monitor) whose mean and contrast were matched to those of the natural scenes. Consistent with the efficient-coding hypothesis they found that, after seconds or minutes of exposure to the natural ensemble, these neurons adapted their response properties so as to increase the information they transmitted.

The appeal of the efficient-coding hypothesis is that it predicts structure in a large and complex data set: the response properties of sensory neurons throughout the nervous system. As such, it is one of the great successes of theoretical neuroscience. But why should the nervous system encode sensory stimuli efficiently? One possible explanation is that neuronal connectivity requires space, and that therefore information must be transmitted using as few 'wires' (axons) as possible[4]. Alternatively, the limiting resource may be the energy associated with neuronal activity[5]; a sparser code using fewer spikes uses less energy.

Finally, the motivation might (as Barlow[1] supposed) be computational: sparse encoding requires uncovering as much of the underlying structure of the signals as possible. That the nervous system seems to be able to do this is impressive, because it is not an easy task. As anyone who has ever tried to write a *Nature* letter knows, writing succinctly requires a very clear idea of what one is trying to say. Or as Mark Twain put it, "I'm sorry this letter is so long, but I did not have time to make it shorter." ■

Michael R. DeWeese and Anthony Zador are in the Cold Spring Harbor Laboratory, 1 Bungtown Road, Cold Spring Harbor, New York 11724, USA.
e-mails: deweese@phage.cshl.org; zador@cshl.org

1. Barlow, H. B. in *Information Processing in the Nervous System* (ed. Leibovic, K. N.) 209–230 (Springer, New York, 1969).
2. Smith, E. C. & Lewicki, M. S. *Nature* **439,** 978–982 (2006).
3. Sharpee, T. O. *et al. Nature* **439,** 936–942 (2006).
4. Chklovskii, D. B., Schikorski, T. & Stevens, C. F. *Neuron* **34,** 341–347 (2002).
5. Laughlin, S. B., de Ruyter van Steveninck, R. R. & Anderson, J. C. *Nature Neurosci.* **1,** 36–41 (1998).

PARTICLE PHYSICS

Quarks on a gravitational string

Nick Evans

Quantum chromodynamics, the theory of the strong nuclear force, is notoriously intractable. An alternative approach brings gravity to bear, and produces fairly accurate predictions of some physical quantities.

The strong nuclear force is the force that causes quarks to bind together to form composite particles, such as the proton. It is explained within the standard model of particle physics by a theory known as quantum chromodynamics (QCD) in terms of fields analogous to electric fields that arise between particles that possess 'colour' charge — the strong-force equivalent of electric charge. Unfortunately for the theorists, however, QCD has consistently eluded analytical solution. The best available calculations rely on huge supercomputer simulations, and the parameters that emerge must be fitted to experimental results.

That situation might change with a remarkable complementary description of the strong force. The new theory incorporates four spatial directions and a dynamic force of gravity that curves space-time in the spirit of the general theory of relativity. As Erlich *et al.*[1] writing in *Physical Review Letters* and Da Rold and Pomarol[2] writing in *Nuclear Physics B* report, it can already be used to produce reasonably accurate quantitative descriptions of light quark–antiquark pairings (mesons) such as the pion.

The genesis of the models is string theory, in which fundamental particles are not points in space, but have an intrinsic length. Seen from afar, the length of these one-dimensional 'strings' is too small to be discerned, but oscillations along this length determine the mass and spin of the particle the string represents. String theory has become a leading contender for a theory of everything, as it can accommodate both gravity and the other forces of nature: so-called open strings have free ends and look like the photon or gluon (the particles that mediate the electromagnetic and strong nuclear force as quanta of the respective fields), whereas closed loops of string have the properties of a graviton, the as yet undiscovered particle that would mediate gravity in a quantum version of general relativity. Problematically, however, string theory turns out to make sense only in a world with nine spatial dimensions, rather than the three that we see.

String theory also contains objects of dimension larger than one: two-dimensional 'membranes', and extensions of still higher dimension known as branes (Fig. 1). These branes emerge in the theory as sub-spaces to which the motion of the open strings, those relevant to QCD, is restricted. Much entertainment has been derived from constructing theories with different matter fields and forces

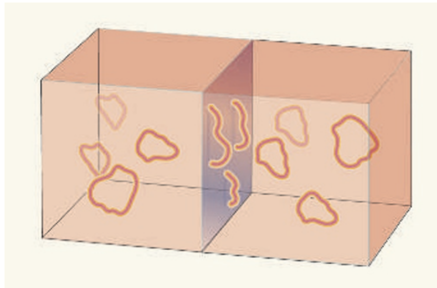

Figure 1 | Braney stuff. In the string theory construction of quantum chromodynamics (QCD), open strings, which describe quarks, are restricted to a membrane-like surface, known as a brane. Closed strings live in a higher-dimension space off the brane and generate an alternative description of QCD.

restricted to the volumes of branes of different dimensions — including three.

In brane theories, the closed strings that describe gravity continue to move in all nine spatial dimensions. Away from the branes' surfaces, the branes' energy generates curvature in space-time, as does gravity in general relativity. Naively, there are therefore two sectors to brane theories: off the brane, describing a theory with similarities to gravity, and on it, describing a theory such as QCD.

But what if these two sectors are actually two alternative descriptions of the same physics? This was originally conjectured[3] for an idealized version of QCD incorporating constraints on its solution to make the theory easier to handle mathematically. But subsequent work that has aimed to remove these devices seems also to support the supposition.

To appreciate how two different theories can describe the same physics, we must first understand that the behaviour of QCD is very different at different energy scales. Colour charge depends, for instance, on the energy exchanged between two quarks in an interaction. If the energy is large, the coupling is weak; if it is small, the coupling becomes very strong. So, the farther quarks move away from each other (and the less energy is exchanged between then), the stronger they bind. This 'quark confinement' is the reason why we never see a free quark, only those bound into composite particles.

In the gravitational, string-theory description of this phenomenon, the quark coupling is represented by a field that describes the distribution of strings in space. The role of energy scale is assumed by the direction radial from

©2006 Nature Publishing Group