the behaviour is strengthened[11,12]. This mechanism is adaptive in evolutionary terms as it normally causes a fairly selective enhancement of those responses that generate the reward.

The authors hypothesize that the formation of associations between a reward and the representation of elements of a rat's trajectory in the immediate past is boosted during sharp-wave-associated replay by a neuromodulatory signal such as dopamine. Dopamine is a chemical released in the forebrain (in the striatum and cortex, and presumably the hippocampus) at the time of reward, especially when reward is not expected by the animal[13–15]. Because ripple trains are variable in length, the effects of the boosting signal would be most reliable if it occurred at the beginning of the sharp wave; however, an early boost could be linked to the key later elements of the preceding firing sequence only if the sequence were reactivated in reverse order, as in Foster and Wilson's study. It remains to be seen whether these speculations will stand up to experimental testing. At the moment, we do not know whether dopamine-releasing neurons fire in synchrony with hippocampal sharp waves.

If reverse replay is a mechanism for strengthening hippocampal sequence memories during goal-directed behaviour, several questions arise. For example, is the firing sequence stored as an ordered memory or as a unitary representation with a stronger representation of the later than of the earlier elements? Moreover, is reverse replay specific to sharp waves that coincide with reward? Sharp waves are observed during breaks without rewards. Do these sharp waves also exhibit reverse replay and, if so, are these associated with memory storage? Finally, can memories of events be stored without interleaved sharp waves? Whatever the answers may be, the discovery of reverse replay is bound to pave the way for more surprises. ■

Laura L. Colgin and Edvard I. Moser are at the Centre for the Biology of Memory, Norwegian University of Science and Technology, MTFS, 7489 Trondheim, Norway.
e-mail: edvard.moser@cbm.ntnu.no

1. Stickgold, R. *Nature* **437,** 1272–1278 (2005).
2. Foster, D. J. & Wilson, M. A. *Nature* **440,** 680–683 (2006).
3. Squire, L. R., Stark, C. E. & Clark, R. E. *Annu. Rev. Neurosci.* **27,** 279–306 (2004).
4. Wilson, M. A. & McNaughton, B. L. *Science* **265,** 676–679 (1994).
5. Lee, A. K. & Wilson, M. A. *Neuron* **36,** 1183–1194 (2002).
6. Buzsaki, G., Leung, L. W. & Vanderwolf, C. H. *Brain Res.* **287,** 139–171 (1983).
7. Klausberger, T. *et al. Nature* **421,** 844–848 (2003).
8. Buzsaki, G. *Neuroscience* **31,** 551–570 (1989).
9. O'Neill, J., Senior, T. & Csicsvari, J. *Neuron* **49,** 143–155 (2006).
10. O'Keefe, J. & Dostrovsky, J. *Brain Res.* **34,** 171–175 (1971).
11. Skinner, B. F. *The Behavior of Organisms* (Appleton-Century-Crofts, New York, 1938).
12. Hull, C. L. *A Behavior System* (Wiley, New York, 1952).
13. Berridge, K. C. & Robinson, T. E. *Brain Res. Rev.* **28,** 309–369 (1998).
14. Montague, P. R., Hyman, S. E. & Cohen, J. D. *Nature* **431,** 760–767 (2004).
15. Schultz, W. *Annu. Rev. Psychol.* **57,** 87–115 (2006).

QUANTUM METROLOGY

Size isn't everything

Samuel L. Braunstein

From probing living cells under a microscope to scanning the heavens for gravity waves, the limitations of precision measurements constrain our capacity to discover more about the world. But what exactly are those limits?

Just how accurate can measurements get? Whereas classical physics places no fundamental limits on how well we can do, in the quantum world it's a different story. Writing in *Physical Review Letters*, Giovannetti, Lloyd and Maccone[1] derive general limits for the precision with which a single variable can be measured quantum mechanically.

But is this new? After all, Heisenberg's uncertainty principle — one of the earliest results in quantum mechanics — already places a fundamental limitation on the precision with which we can make a measurement. In its simplest form, the uncertainty principle identifies so-called complementary observables, pairs of quantities for which knowing one quantity precisely means that the other can only be poorly known. This fundamental principle makes it impossible to learn everything about a quantum-mechanical system.

If we monitor only one quantity, however, there is no such in-principle limitation. In fact, this is exactly the strategy exploited in interferometric measurements, in which light travels down a pair of distinct paths and the difference between the two path lengths leads to an observable change in the output of the device. This path difference can be measured to an arbitrary accuracy. But what if we are given some constraint, such as a total energy budget or total light intensity? We all know that it is easier to see in a well-lit room than in a dim one. Similarly, the higher the energy or light intensity in an interferometer, the higher its resolution. One may therefore ask, for a fixed budget, how small a path difference can be discerned?

Our intuition from everyday experience tells us that the most promising strategy for measuring a distance is to choose a measuring stick with marked intervals of length comparable to the distance we wish to measure. We would not, for example, choose a metre stick to measure a molecule. Following similar logic, we might choose the wavelength of light for our interferometer to be comparable to the path difference we want to measure. Surprisingly, Giovannetti and colleagues' latest result[1] can be used to show that, for optimal quantum strategies, there is no such bias to the size of our measuring stick or the separation of its tick marks.

An optimal strategy refers to a measurement procedure that minimizes the effects of noise on a signal. Ultimately, any measurement is limited by the amount of noise in the system: to discern a signal, the signal-to-noise ratio should be around one or larger. This premise underpins all parameter-estimation theory, both classical and quantum. Classically, statistical averaging over N repeated but independent measurements will lead to a $\sqrt{N}$ reduction in the noise. This improvement is known to be optimal because it achieves the bound, known as the Cramér–Rao lower bound[2], that expresses the best accuracy that can be accomplished in the statistical estimation of a parameter. When this classical bound is generalized to repeated quantum measurements, the analogous quantum bound provides a tighter form of the uncertainty principle recast in the language of parameter estimation[3]. However, quantum theory allows much more freedom in choosing measurement strategies than is possible in the classical world.

One of the most bizarre features of the quantum world is quantum 'entanglement', which allows systems to exhibit stronger correlations than are possible classically. Using entanglement and other tricks, quantum mechanics has led us to devise sophisticated information-processing algorithms that one day may lie at the heart of the enormous speed-ups promised by quantum computation. For example, searching for a needle in a haystack would be much faster — in principle — on a quantum computer than a classical one. The possibility of using entangled systems and/or entangled measurements, and sophisticated algorithms built into measurement devices, raises questions about the ultimate (most general) quantum bounds to measurement.

Giovannetti and colleagues' key insight[1] into this question is to recast the measurement process in terms of quantum circuits, analogous to electrical circuits, with various quantum gates, similar to logic gates, representing different quantum-mechanical 'operators'. They then introduce black-box operators that perturb the quantum state in a known fashion, but by an unknown amount. Such an operation might, for instance, be adding a phase delay along one arm of an interferometer: the unknown parameter associated with the black box thus corresponds to the parameter we

would like to estimate. Once such a black box is conceptualized, it may be reused in the circuit again and again (each black box having the same unknown parameter). The beauty of this language lies in its generality, which allows a rich class of measurement strategies involving N such identical black boxes in a circuit of arbitrary design.

Using this formalism, Giovannetti *et al.* show that the optimal accuracy achievable in estimating the value of the black-box parameter can be obtained in a simple circuit with N black boxes, running on an N-fold entangled state. Surprisingly, recourse to entangled measurements (joint measurements of multiple paths of the circuit), or rearrangements of the circuit to correspond to sophisticated quantum-search strategies, will not lead to any further improvement.

What is this optimal performance? In fact, it depends entirely on the range of observable values of the black-box operator. In any circuit with N black boxes, the noise associated with the estimation of the black boxes' parameter will be reduced at most N-fold compared with the noise in the best circuit with only a single black box. That represents a considerable advantage over the $\sqrt{N}$ improvement of the classical case. The good (and reassuring) news is that this limit is exactly what one would have expected from a naive application of the good old Heisenberg uncertainty principle: it is none other than the Heisenberg limit.

So what relevance does all this have to the choice of size in our metre sticks? Well, let's return to our interferometer. For a given energy budget (or light intensity), but freedom in our choice of wavelength, we would naively expect the shorter wavelength to yield higher sensitivity. However, the longer the wavelength, the more photons we can squeeze into our interferometer. In other words, with the same budget, we can sample the black box exactly that many more times. Indeed, the Heisenberg-limited measurement is equally good, independent of our choice of measuring stick.

Two limitations to the strategy of Giovannetti *et al.*[1] lie in the quantum version of the Cramér–Rao bound on which it is based[3]. First, this bound can be reached only for problems involving single-parameter estimation, so extensions to multiple parameters may lead to different results. For instance, the estimation of the orientation of quantum spins (involving two unknown angles in three-dimensional space) can be enhanced by entangled measurements[4]. Second, the Cramér–Rao bound can be achieved only for an infinite number of repeated measurements. Thus, a result that expresses the approach to this asymptote would fill a gap in our current understanding. Indeed, it may be just this discrepancy that underlies the enhanced precision in determining the orientation of quantum spins using entangled measurements — an enhancement that vanishes in the limit of an infinite number of spins[4].

Currently, we are far from putting the ultimate bounds described by Giovannetti *et al.*[1] into practice. One example would be the Laser Interferometer Gravitational-Wave Observatory (LIGO), an exciting experiment that aims to detect tiny ripples in the fabric of space-time. The LIGO interferometer currently implements only classical strategies scaling as $1/\sqrt{N}$ (where N is the number of photons in the interferometer). In its current set-up, LIGO requires a circulating power of 10–20 kilowatts to achieve minimal sensitivities for detecting gravity waves. In principle, if we could implement a quantum-limited scheme, a similar sensitivity could be achieved with only nanowatts. Such prospects promise an even brighter future for gravity-wave astronomy in the long term — and for precision measurement in general. ■

Samuel L. Braunstein is in the Department of Computer Science, University of York, York YO10 5DD, UK.
e-mail: schmuel@cs.york.ac.uk

1. Giovannetti, V., Lloyd, S. & Maccone, L. *Phys. Rev. Lett.* **96,** 010401 (2006).
2. Cramér, H. *Mathematical Methods of Statistics* 500–504 (Princeton Univ. Press, 1946).
3. Braunstein, S. L. & Caves, C. M. *Phys. Rev. Lett.* **72,** 3439–3442 (1994).
4. Gill, R. & Massar, S. *Phys. Rev. A* **61,** 042312 (2000).

MATERIALS SCIENCE

Nanostructures in a new league

John J. Rehr

Aperiodic materials do not surrender details of their structure as readily as do their crystalline counterparts. The latest computational solution to this problem brings aspects of 'the beautiful game' into play.

Investigations of crystalline materials through X-ray and neutron diffraction have been a triumph of experimental science, allowing structures ranging from complex minerals to proteins and DNA to be unravelled[1]. But how can the structure of a material that is aperiodic — one that is non-crystalline, or cannot be crystallized — be determined? On page 655 of this issue[2], Juhás *et al.* present an intriguing solution to this question with a novel algorithm for reconstructing three-dimensional structures from 'pair distribution function' (PDF) data. Aperiodic materials are among the technologically most interesting nanoscale materials currently under study, and the approach could be widely applicable.

Several techniques exist for determining the local, atomic-scale structure of materials. These range from scanning tunnelling microscopy (STM) to spectroscopic methods that use X-rays, such as extended X-ray absorption fine structure (EXAFS) analysis. Each has its advantages and drawbacks. STM can give beautiful images, although not in three dimensions. For structural information to be inferred from spectroscopic techniques such as EXAFS, an accurate theoretical model relating spectra to structure is required[3].

PDF analysis avoids some of these problems because it solely involves data on the distribution of distances between atoms in a structure — information that is readily obtained from X-ray or neutron-scattering experiments[3]. Why, then, is PDF not the method of choice for structure determination? The first factor is data quality: although the PDF technique has been known for decades, the lack of high-resolution data has limited its applicability, as well as that of many other techniques. That situation is now changing with the latest generation of experiments using modern neutron and synchrotron X-ray sources.

The second crucial factor is that an algorithm must be found that solves the 'inverse problem'; that is, given a set of experimental data, how to extract the three-dimensional structure that must have created it. Determining the structure corresponding to PDF data, the question tackled by Juhás and colleagues[2], is just such an inverse problem. The inverse problem is usually not trivial, as it involves various assumptions about a material and, potentially, many material-dependent parameters. Solutions typically involve minimizing the mean squared deviation between the experimental data and the data predicted from a theoretical model of the structure. This process often needs significant computational resources, as it requires the 'direct problem' — that is, a theoretical model for the experimental signal resulting from a trial structure — to be solved many times in the process of finding the minimum.

Obtaining a solution to the inverse problem is equivalent to an optimization strategy for finding the global minimum of a quantity involving many variables among a forest of possible minima. Numerous advances have been made in such strategies, which are crucial in fields from economics to protein folding[4]. These include the development of 'genetic' algorithms inspired by the rules of evolutionary biology, and 'simulated annealing' techniques that mimic the way metals freeze into a