COGNITIVE NEUROSCIENCE

Decision amid uncertainty

Jonathan D. Cohen and Gary Aston-Jones

Choosing whether to stick to a belief or to abandon it in the face of uncertainty is central to human behaviour. Modelling implicates brain chemicals called neuromodulators in adjudicating this essential decision.

Why did you choose to read this article? Perhaps you are a neuroscientist eager to refine your knowledge. Or perhaps you are keen to broaden your horizons outside your current discipline. These motivations reflect a fundamental trade-off in how we invest our time and effort: individuals must continually decide whether it would be better to pursue known sources of reward, or whether there is more to be gained by searching for new strategies or opportunities. In reinforcement learning, this dilemma is referred to as the trade-off between exploitation and exploration. There is growing evidence that the mechanisms used to resolve this trade-off are directly regulated by neuromodulators[1–3]. Yu and Dayan[4], writing in *Neuron*, extend this work by using simple principles from Bayesian probability theory to derive a sophisticated model of how neuromodulatory systems are central to the trade-off decision.

As their name suggests, neuromodulators such as dopamine, acetylcholine and noradrenaline seem to modify the effects of neurotransmitters — the molecules that allow communication between neurons. Neuromodulatory systems are implicated in almost every mental function, including attention, learning and emotion, and they are disturbed in major neurological and psychiatric disorders, from Alzheimer's disease and post-traumatic stress disorder to depression and schizophrenia.

The conventional view of neuromodulators has been that they have broad, nonspecific functions, such as signalling reward (dopamine) and regulating arousal (noradrenaline). But a current renaissance in the field is showing that neuromodulators have more-specific functions in learning and decision-making. For example, dopamine neurons are implicated in signalling errors in reward prediction, a role that is central to reinforcement learning[1,5]. Moreover, noradrenaline may be key in facilitating the responses to decision-making processes[6,7], and in regulating the balance between exploitation and exploration[3,8].

To examine the role of neuromodulators in such processes, Yu and Dayan[4] exposed subjects to a task involving a set of cues — differently coloured arrows pointing to the left or right — one of which points to where a target will subsequently appear. The participant must respond as quickly as possible to the target; if they work out what the predictive cue is, they tend to do better. In a typical set-up, the predictive cue stays the same for a number of trials, but then changes without warning.

The crux of the task is that there are two forms of uncertainty associated with the cues. Subjects must work out which cue predicts where the target will appear, as well as how reliably it does so. If the predictive cue were 100% accurate, the task would be trivial: most people would quickly discover which arrow reliably points to the target, and notice as soon as this changed. However, the cue is usually set to be only partly predictive (for example, 80% of the time). And herein lies the intrigue — and generality — of the problem. Suppose you have come to believe that a particular cue predicts the target, but in the last few trials it has failed to do so. How do you know whether this is because the cue is not a perfect predictor of the target (like most cues in the world), or because the relevant cue has changed? More generally, how do we decide whether to stand by our beliefs, even as we recognize their fallibility, or to abandon them in search of better ones?

Yu and Dayan cast this dilemma in terms of a distinction between expected uncertainty (in their task, the less-than-perfect reliability of a cue) and unexpected uncertainty (a surreptitious switch in the relevant cue). They propose that information about these forms of uncertainty is coded in the brain by different neuromodulatory systems — with acetylcholine reflecting the degree of expected uncertainty, and noradrenaline gauging unexpected uncertainty.

Yu and Dayan develop a model of how acetylcholine and noradrenaline levels encode uncertainty, and how their interaction determines whether we should abide by or abandon an existing belief. Their analysis implies that if the optimal strategy was computed on each trial, the process would be so demanding as to be biologically unfeasible. However, they demonstrate that an alternative probability algorithm that is biologically plausible can approximate the optimal strategy.

Their model allows Yu and Dayan to make detailed quantitative — and sometimes counterintuitive — predictions about neuromodulatory function and its influence on behaviour. For example, according to their theory the degree of unexpected uncertainty that causes you to abandon a belief should depend on the level of expected uncertainty; that is, if you know that the selected cue is not reliable, you will have a higher tolerance for its failure to predict the target. This makes interesting predictions about how disturbances of acetylcholine and noradrenaline levels will affect behaviour. Confirmation of these predictions, in turn, is likely to provide deeper insight into the patterns of behavioural deficits observed in clinical disorders involving disturbances

---

MATERIALS SCIENCE

Sticky business

Geckos are small tropical lizards known for their ability to attach themselves to walls and ceilings. Their sticking power has been attributed to the millions of little hairs on their feet: each hair exerts a tiny force on the surface, which, added together, enables geckos to hang off surfaces at any angle — a phenomenon that has been exploited to produce adhesive 'gecko tape'.

But what happens to water droplets on a surface that has the potential sticking qualities of a gecko's foot but that is also extremely hydrophobic? M. Jin and colleagues produced such a superhydrophobic, adhesive surface by packing polystyrene nanotubes into an array at a density of 6 million per square millimetre (*Adv. Mater.* doi:10.1002/adma.200401726). The water contact angle — a measure of surface wettability — for the nanotube array was 162°, compared with 95° for a smooth polystyrene surface.

Using a microelectromechanical balance, Jin and colleagues found that ten times more force was required to remove water from the nanotube surface than from a normal superhydrophobic surface. They attribute this increase to the same mechanism — cumulative van der Waals forces — that governs gecko adhesion. They also discovered that water droplets weighing up to 8 milligrams stuck to the nanotube surface regardless of the angle at which the surface was held (see images), and that the adhesive strength of the surface increased with increasing density of the nanotubes.

Water droplets on the superhydrophobic structured surface could be transferred to a hydrophilic surface without any loss of fluid or introduction of contamination — properties that the authors hope might be useful for practical fluid manipulation.

Rosamund Daw

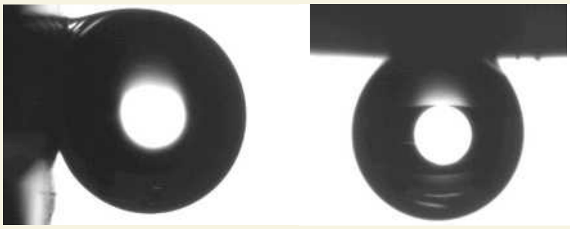

WILEY

of attention, decision-making and learning.

It is perhaps in this regard that Yu and Dayan's work is most noteworthy. Research on the role of neuromodulatory disturbances in mental disorders has tended to focus on simple hypotheses concerning static excesses or deficits of activity in individual neuromodulatory systems, with little consideration of interactions between systems. A more profound understanding of these dynamics, and their relationship to cognition and behaviour, is crucial if we are to understand how disruption of these systems contributes to the clinical symptoms associated with neurological and psychiatric disorders, and ultimately, how to design effective treatments.

A future problem will be to bring this theory into contact with ones that have framed the question more specifically in terms of the trade-off between exploration and exploitation, and the maximization of reward (for instance, theories about reinforcement learning from neuroscience and utility maximization from economics). The theory also needs to address the high temporal specificity of neuromodulatory systems, which can show rapid phasic responses following task-relevant events (within 100–200 ms), and which may have an immediate impact on task performance[2,6,9–11]. Finally, an understanding of the biophysical and circuit mechanisms by which acetylcholine and noradrenaline interact to produce the proposed functions is needed to advance this view of neuromodulators in brain function.

Yu and Dayan's work is an impressive contribution to the evidence that neuromodulatory function is more specific than previously thought. We can say with some certainty that their theory represents a particularly promising direction of research, and that their paper is a highly rewarding read. ■

Jonathan D. Cohen is in the Department of Psychology and the Center for the Study of Brain, Mind and Behavior, Green Hall, Princeton University, Princeton, New Jersey 08544-1010, USA. Gary Aston-Jones is in the Laboratory of Neuromodulation and Behavior, Department of Psychiatry, University of Pennsylvania School of Medicine, 705 Stellar Chance, Philadelphia, Pennsylvania 19104-6100, USA.
e-mail: gaj@mail.med.upenn.edu

1. Schultz, W., Dayan, P. & Montague, P. R. *Science* **275,** 1593–1599 (1997).
2. Hasselmo, M. E. & McGaughy, J. *Prog. Brain Res.* **145,** 207–231 (2004).
3. Aston-Jones, G. & Cohen, J. D. *Annu. Rev. Neurosci.* **28,** 403–450 (2005).
4. Yu, A. J. & Dayan, P. *Neuron* **46,** 681–692 (2005).
5. Montague, P. R., Dayan, P. & Sejnowski, T. J. *J. Neurosci.* **16,** 1936–1947 (1996).
6. Clayton, E. C., Rajkowski, J., Cohen, J. D. & Aston-Jones, G. *J. Neurosci.* **24,** 9914–9920 (2004).
7. Brown, E. T. *et al. Int. J. Bifurcat. Chaos* **15,** 803–826 (2005).
8. Usher, M., Cohen, J. D., Servan-Schreiber, D., Rajkowski, J. & Aston-Jones, G. *Science* **283,** 549–554 (1999).
9. Hasselmo, M. E. *Behav. Brain Res.* **67,** 1–27 (1995).
10. Aston-Jones, G., Rajkowski, J., Kubiak, P. & Alexinsky, T. *J. Neurosci.* **14,** 4467–4480 (1994).
11. Hasselmo, M. E. *Neuron* **46,** 526–528 (2005).

QUANTUM OPTICS

Crystal-clear images

Claude Fabre

Two experiments that use nonlinear crystals to control the spatial distribution of photons in optical images bring the field of quantum imaging closer to maturity. Quantum information processing could ultimately benefit.

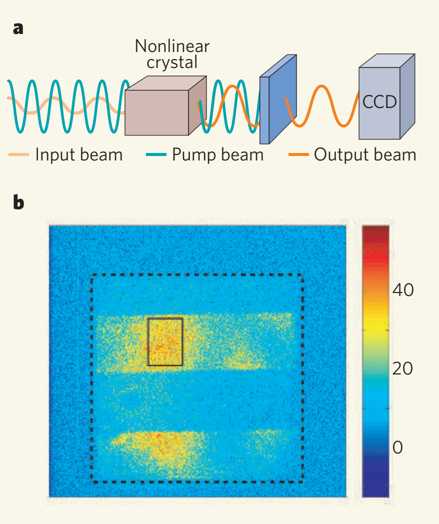

Figure 1 | Parametric amplification of light. **a,** Principle of a parametric optical amplifier. The nonlinear crystal allows energy transfer between the two optical beams that propagate through it: the intense pump beam (green) oscillating at frequency f and the weak input beam (orange) oscillating at frequency $f/2$. When the input beam has a well-defined phase relation with the pump beam, it is amplified when it propagates in the crystal. This amplification process is 'noiseless' — no extra quantum fluctuations are added to the input signal. The filter behind the crystal blocks out the pump beam, and the output is recorded on a charge-coupled device, or CCD. **b,** An amplified image recorded on the CCD when the input image is a single rectangular slit of constant intensity (lighter area at the centre). Strong spatial fluctuations between pixels are observed in the output image, related to the randomness of the impact positions of the photons making up the image. By measuring the spatial intensity averages in the rectangular region inside the solid black line (3,266 pixels), Mosset and colleagues[1] show that the amplification process does not degrade the signal-to-noise ratio of the image. The dashed square shows the projection of edges of the crystal.

God, it seems, plays dice. Randomness lies at the root of many quantum effects: in the arrival times of the photons on a detector, for example; or, if the intensity is such that the photons cannot be distinguished individually, in the temporal fluctuations of the current generated in the detector by those photons. In the past two decades, physicists have found ways to tame this 'quantum noise', and to master, at least in some instances, the temporal distribution of photons.

Two papers now show that it is also possible to control photons' spatial distribution. Writing in *Physical Review Letters*, Mosset and colleagues[1] provide the first experimental demonstration of an optical device that amplifies an image without increasing the spatial distribution of where the photons hit. In a contribution to the same journal last year, Jedrkiewicz *et al.*[2] used a similar device to produce two images with photons that are identically distributed in space.

Because of the quantum nature of light, the amplification of an optical signal is not a simple multiplication of the input signal by the gain. For very large gains, the signal-to-noise ratio at the output is a factor of two smaller than that at the input[3]. Fortunately, a special class of amplifiers that is not doomed to degrade the signal being amplified does exist. The 'degenerate parametric amplifier' is one of these (Fig. 1a). This amplifier consists of a crystal that responds in a nonlinear manner when it is subjected to an intense 'pump' beam, the electric field E of which varies cosinusoidally with time t as $E = E_0\cos(\pi ft)$. (E_0 is the maximum amplitude of the wave, and f is the frequency of its oscillation.) An optical input signal at a frequency $f/2$ can be amplified while propagating in this crystal without changing the signal-to-noise ratio — that is, without degrading the quality of the information channel[3]. The drawback of this type of amplifier compared with a more conventional one is that it is 'phase-sensitive': only cosinusoidally varying fields of the form $E = E_0\cos(\pi ft)$ are amplified, whereas sinusoidally varying fields $E = E_0\sin(\pi ft)$ are attenuated.

The first experimental demonstration of this 'noiseless amplification' was made in 1993 for single-channel signals (those that have a single frequency and a fixed transverse shape), such as are used in optical telecommunications[4], and in 1999 for signals of an arbitrary transverse shape — images[5].

The novel aspect of the work of Mosset *et al.*[1] is its emphasis on spatial, rather than temporal, quantum noise. Spatial quantum fluctuations appear only in images, not in single-channel signals. They are related to the degree of randomness that exists in the distribution of photons composing a beam of light carrying an image; this distribution is measured in the plane transverse to the direction of propagation. In essence, the authors show that images are not degraded by additional spatial quantum fluctuations when they are amplified by a degenerate parametric amplifier.