WEB WATCH

Sweet dreams

In Greek mythology, Hypnos is the God of Sleep. He fathered many sons, but only three of them were chosen to rule the realm of dreams — Phantasos, who induces dreams of inanimate objects, Phobetor, who elicits images of animals, and Morpheus, the true God of Dreams, who populates our nights with images of humans. This ancient fascination with sleep and dreaming is still present in our times, and we continue to learn more and more about both phenomena as the tools of modern biology bring their strengths to bear on the problem. If you want to get a clear idea of how much we already know about sleep, then you must pay a visit to Basics of Sleep Behavior, a fantastic site where you can get up to speed on what happens in our brain while we are “in the arms of Morpheus”.

Basics of Sleep Behavior is an online tutorial developed by the Sleep Research Society and the Brain Information Service at the University of California, Los Angeles. It includes accessible discussions of different aspects of sleep physiology and pathology that are aimed at a wide audience — from undergraduates to established researchers who might want to learn the basics of the field. It also includes a helpful bibliography (which would benefit from some updating), a dictionary, and links to another very useful resource — the Sleep Home Pages.

The Sleep Home Pages is a portal, also maintained by the Brain Information Service, which provides an almost inexhaustible list of links to numerous organizations, discussion forums, publications and all kinds of pages related to sleep. Together, Basics of Sleep Behavior and the Sleep Home Pages are ideal resources to start exploring the domains of Hypnos and his cohort of children.

Juan Carlos López

LEARNING THEORY

Trial and error

Our elders and betters often counsel us to learn from our mistakes. Learning theory states that we learn that a stimulus is paired with a reward only if, initially, we don't expect it — there must be a ‘prediction error’ for learning to occur. Now Waelti *et al.* have shown that dopamine neurons in the midbrain seem to follow the same rules. Their study is a compelling example of the direct testing of a prediction that arises from associative learning.

Classical learning theory predicts that learning will occur whenever a stimulus is paired with a reward, and this seems intuitively sensible. But more recent work has led to the idea that learning only occurs when a prediction error is present. This can be shown using a ‘blocking’ procedure. First, an animal learns, by repeated trials, that stimulus A — a bell, perhaps — is always followed by a reward — say, some fruit juice. After a while, the animal will lick at the juice spout every time it hears the bell, in anticipation of the juice. If the animal then sees a coloured light (stimulus X) together with the bell before the juice is delivered, we might expect that it would learn to associate the light with the juice, and lick the spout even if it saw the light alone. But this does not happen — because the juice is already fully predicted by the bell, there is no prediction error, so the animal never learns to associate the light and the juice.

Waelti *et al.* have recorded the activity of dopamine neurons in the midbrain during this type of training, with various coloured shapes as stimuli. Dopamine signalling is thought to be important for reward systems, and there is evidence that the activity of dopamine neurons may code the prediction error.

In fact, in the new study, Waelti *et al.* found that dopamine neurons were activated by the reward-predicting stimulus A, but not by a stimulus that was not paired with reward. When the authors trained their animals using a blocking procedure, they found that the compound stimulus set (AX) activated the dopamine neurons, but X alone did not.

By contrast, if another stimulus, B — perhaps a different coloured light — is not normally paired with reward, but then B and Y (a whistle) are together paired with the juice, the animal will learn that Y predicts juice even in the absence of B. This is because B does not predict any reward, so when the two stimuli are paired with the juice, there is a prediction error, which leads to learning. After training, dopamine neurons were strongly activated by stimulus Y.

The neuronal responses precisely reflected the behavioural responses of the animals, which showed much more anticipatory licking following Y than after X. When a reward was presented after stimulus X, it strongly activated the dopamine neurons — exactly as we would expect if the neurons code prediction error. By contrast, after stimulus Y, which the animals had learned to associate with reward, the presentation of

fruit juice would create no prediction error and, correspondingly, it produced no increase in dopamine activity.

The dopamine neurons appear to ‘learn’ the stimulus–reward association, and their responses conform precisely to the predictions of learning theory. Neuronal learning requires the presence of a prediction error, and the neuronal responses appear to code this prediction error. It is possible that similar approaches that aim to integrate electrophysiological recording and behavioural learning rules will lead to important insights into the cellular basis of learning.

Rachel Jones

References and links

ORIGINAL RESEARCH PAPER Waelti, P. *et al.* Dopamine responses comply with basic assumptions of formal learning theory. *Nature* **412**, 42–48 (2001)

FURTHER READING Schultz, W. & Dickinson, A. Neuronal coding of prediction errors. *Annu. Rev. Neurosci.* **23**, 472–500 (2000) | Schultz, W. Multiple reward signals in the brain. *Nature Rev. Neurosci.* **1**, 199–207 (2000)

WEB SITE Schultz's lab