Neuroscience

Speed of learning depends on turning

István Taisz & Gregory S. X. E. Jefferis

The neurotransmitter dopamine has been shown to serve as a signal for learning in the fly's navigation centre. The rate at which the fly learns depends on turning, so only useful visual information is used to update the fly's mental map. **See p.316**

Nervous systems face a trade-off between storing and updating existing representations of the world. Adaptively controlling a system's plasticity could help to find the balance, allowing fresh information to be learnt only when it is relevant. But what could serve as the 'when-to-learn' signal? On page 316, Fisher *et al.*[1] describe a solution to this question for the neural network that encodes head direction in fruit flies. The authors link changes in learning rate to turning-driven release of the neurotransmitter molecule dopamine.

To find a destination, it is useful to keep track of which way you are facing. When fruit flies navigate towards a banana or an egg-laying site, they learn a representation of their orientation using visual landmarks and cues about their own motion[2]. This 'heading estimate' can then be matched with the direction taken to reach a desired stimulus, or be used to estimate the fly's own velocity, thereby allowing the insect to memorize its route.

At the core of this system are EPG neurons, which receive information about landmarks from visual neurons and store the fly's heading estimate. EPG neurons form a ring-shaped structure in the brain, called the ellipsoid body. A local increase in the activity of neurons in one portion of the ring encodes the fly's heading angle; as the fly turns, activity travels around the ring to update the angle. The exact mapping between the world and the fly's compass is arbitrary: it is not that the same portion of the ring is active every time a fly faces north, for instance. Instead, one random but stable map emerges for any given visual scene. Every time a fly encounters that particular scene, the same location on the ring is activated when it faces a certain point.

This combination of flexibility and stability is possible because the synaptic connections from visual neurons to EPG neurons are plastic: their connection strength can change rapidly to reconfigure the heading map[3,4]. One could imagine that such a system, which continuously encodes spatial information, might also continuously learn the structure of the environment. Alternatively, learning could be limited to moments when signals convey the richest information about the visual scene.

Fisher *et al.* probe these alternatives by asking whether and when a group of dopamine-releasing neurons called ExR2 neurons can change the rate at which connection strengths (synaptic weights) between visual and EPG neurons are updated in the ellipsoid body (Fig. 1). In reinforcement learning, dopamine is released when the outcome of our actions does not match what the brain predicts. This compresses incoming data, updating our world model only when it is incorrect. However, representing a heading is a different type of learning: the brain updates its model of the world without a prediction error. What, if anything, is the role of dopamine in this case?

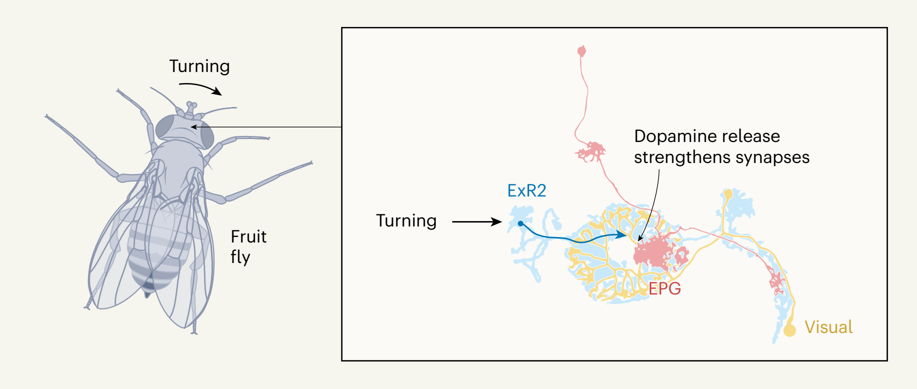

Figure 1 | The neurotransmitter dopamine tells flies when to learn. EPG neurons act as a compass that enables flies to learn their heading (which way they are facing in a scene). Fisher *et al.*[1] have found that neuronal signals produced when a fruit fly turns activate dopamine-producing ExR2 neurons. The resulting dopamine release increases the strength of connections called synapses between visual neurons and EPG neurons, and so the rate of learning in EPG cells. In this way, only information-rich visual data are incorporated in the fly's mental estimate of heading. (Neurons are depicted as an anatomically accurate map.)

To answer this question, the authors engineered fruit flies such that the ExR2 neurons expressed a protein called a genetically encoded calcium indicator, which gives a readout of neuronal activity. They recorded ExR2 activity while the flies walked on an air-suspended ball in the dark (meaning that the insects had no visual input). The authors found that ExR2 activity closely followed moment-to-moment changes in the flies' rotational speed: the faster the flies turned around, the greater the activity in these neurons. Incorporating these findings into an earlier model[4], the authors showed that dopamine release, driven by rotational signals, modulated learning rates to produce a precise representation of heading. This suggests that learning is discontinuous. A fly's turning behaviour acts as a signal that richer visual information should now be available.

To test experimentally whether dopamine controls changes in synaptic weights in this system, the authors measured the influence of visual signals on EPG neurons while manipulating the activity of ExR2 neurons. This time, they chose an approach called patch-clamp electrophysiology to record the electrical activity of single EPG neurons, and expressed an ATP-gated cation channel protein in ExR2 neurons – a modification that allowed them to control the neurons' dopamine release. In a virtual-reality set-up, a bar rotated around a fly that was fixed in position. This visual cue caused inconsistent responses in EPG neurons. After briefly stimulating ExR2 neurons, the responses became more stably related to the stimulus, demonstrating that ExR2 neurons strengthen the influence of visual information. Next, by imaging the activity of the EPG population, the authors confirmed that the entire head-direction map showed strengthened association with a visual cue after ExR2 activation.

Finally, to show that the ExR2 neurons are essential for creating an accurate representation of heading, the authors dampened the activity of these neurons by engineering flies to express a particular potassium channel protein. In this experiment, the group again placed flies in a virtual-reality system; but this time, the animals' movement caused a visual stimulus to move around them. In control animals, this established a strong coupling between cue direction and heading estimate in EPG neurons; this connection was maintained for minutes, even after the authors uncoupled the stimulus position from the fly's turns. But dampening ExR2 neurons decreased the strength of the coupling – showing that this cell type is necessary for generating an accurate heading representation based on a visual landmark.

The role of dopamine in reinforcement learning is long established. But now, Fisher *et al.* propose that the molecule also controls learning rate in unsupervised learning; in this case, a behavioural variable (the rate of turning) signals, through dopamine release, to compress the flood of incoming sensory information, allowing only the most useful to be selected. The authors point out that when a fly walks in a straight line, its visual information is limited, so a mental model built on such partial data is likely to be incorrect. This means that coupling the update rate in the compass with turning is not only efficient – similar to updating only those pixels that differ between frames when streaming a video – but also avoids the risk of overfitting to incomplete information.

Several questions remain, including how exactly dopamine changes synaptic connection strength. For example, does this occur only through changes in the number of neurotransmitter receptors, or do other structural changes also contribute to learning? It is unclear how general the when-to-learn signal is. The arthropod navigation centre is widely conserved in evolution, but locomotion strategies are diverse: how do circuits adapt to these differences? In mammals, the same task is performed by a functionally analogous, but structurally distinct, system that might use other motor signals and molecules to create its own learning logic. The mammalian navigation system not only stores spatial maps, but can also represent more-abstract cognitive spaces[5,6]. How does the brain find an appropriate when-to-learn signal for these maps? Answers could inspire the design of artificial-learning systems that can select task-relevant data on a moment-to-moment basis.

István Taisz & Gregory S. X. E. Jefferis are in the Division of Neurobiology, MRC Laboratory of Molecular Biology, Cambridge CB2 0QH, UK.
e-mails: itaisz@mrc-lmb.cam.ac.uk; jefferis@mrc-lmb.cam.ac.uk

1. Fisher, Y. E., Marquis, M., D'Alessandro, I. & Wilson, R. I. *Nature* **612**, 316–322 (2022).
2. Seelig, J. D. & Jayaraman, V. *Nature* **521**, 186–191 (2015).
3. Fisher, Y. E., Lu, J., D'Alessandro, I. & Wilson, R. I. *Nature* **576**, 121–125 (2019).
4. Kim, S. S., Hermundstad, A. M., Romani, S., Abbott, L. F. & Jayaraman, V. *Nature* **576**, 126–131 (2019).
5. Constantinescu, A. O., O'Reilly, J. X. & Behrens, T. E. J. *Science* **352**, 1464–1468 (2016).
6. Nieh, E. H. *et al.* *Nature* **595**, 80–84 (2021).

The authors declare no competing interests.

Quantum information

Driven quantum bits push computational limit

Tomaž Prosen

An exercise in benchmarking a quantum computer reveals that the processor can go beyond the 'integrability' limit, at which dynamical systems no longer have explicit solutions, and standard mathematical techniques struggle. **See p.240**

The old adage that three is a crowd is an apt introduction to the concept of integrability. For example, the celestial mechanics of a star and a planet are said to be integrable. In simple terms, this means that the behaviour of this two-body system can be fully characterized by fewer variables than we have at our disposal: enough of its physical observables remain unchanged for long enough to enable a robust description. But add a third body into the mix (or multiple extra bodies) and this integrability disappears. On page 240, Morvan *et al.*[1] describe how they simulated a model for many-body quantum dynamics in a system of 24 quantum bits (qubits) that are periodically driven to interact with each other. In doing so, the researchers have shown that it is possible to push the limits of integrability.

Integrability applies to phenomena as diverse as the spinning of a top and the propagation of waves in shallow water. It is universally characterized by robust, indestructible excitations that persist for arbitrarily long periods. Such excitations are usually referred to as solitons or quasiparticles, but they can also take the form of stable, bound states that comprise more than one quasiparticle. A system with a large number of components (a many-body system) is said to be integrable if it has an equally large number of physical observables that do not change with time[2,3].

Another physical principle, ergodicity, stipulates that large systems find equilibrium by changing states for long enough to eventually sample all possible states. Integrable systems have unchanging observables, and are therefore not ergodic. For this reason, integrable systems preserve a memory of their initial conditions, and this memory grows with their size. A pioneering example of this phenomenon was demonstrated in an experiment that used ultracold atoms to mimic a Newton's cradle – the

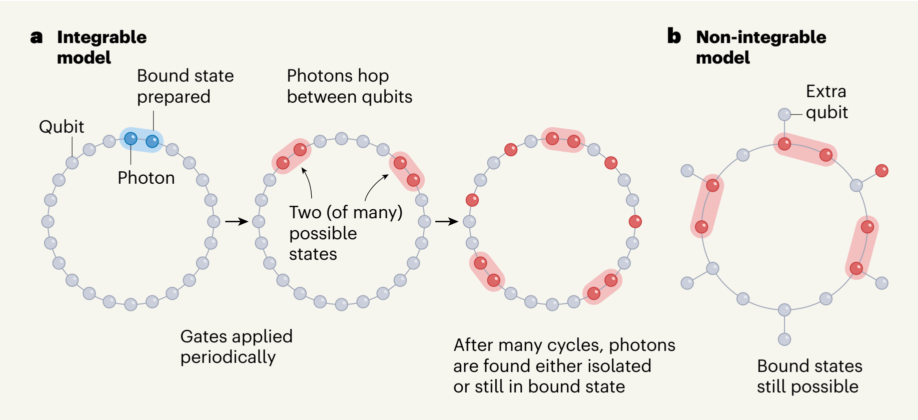

Figure 1 | Driving interactions in a ring of quantum bits. Morvan *et al.*[1] benchmarked a quantum processor with a model that is 'integrable', in that its physical observables permit a robust mathematical description. **a**, The authors prepared a ring of 24 quantum bits (qubits) with a string of 2 or more photons, known as a bound state. They then drove the system by periodically applying a quantum logic gate to all adjacent pairs of qubits in the ring, allowing the photons to move between qubits. With each cycle, single photons hopped around the ring, but so did bound states, as previously predicted for this model[7]. **b**, Morvan *et al.* then implemented a 'non-integrable' model by adding an extra qubit to every other qubit in a smaller system. After many cycles, they found that the bound states persisted for longer than expected. These states should decay to equilibrium quickly in systems that are not integrable, suggesting that quantum processors can push the limits of integrability.