Computational neuroscience

AI networks reveal how flies find a mate

Pavan Ramdya

Artificial neural networks that model the visual system of a male fruit fly can accurately predict the insect's behaviour in response to seeing a potential mate — paving the way for the building of more complex models of brain circuits. See p.1100

From the perspective of a male fruit fly, rotting fruit represents an opportunity. There, he can find and court a potential mate by chasing her, extending his wings and vibrating them to produce a fine-tuned 'love' song. Importantly, this courtship ritual depends on the male suitor receiving visual feedback from the female of interest[1] — but how does the male's eye instruct his brain during this courtship dance? Across species, visual systems transform patterns of light into features that are meaningful to an animal's behaviour: for example, a looming shadow can trigger certain groups of neurons known as feature detectors, which drive fast and reliable predator avoidance[2]. Feature detectors can also be combined to discern more complex visual patterns. On page 1100, Cowley *et al.*[3] combine a machine-learning tool called an artificial neural network with genetics to explore how the male fly uses feature detectors to coordinate its movements with those of its dance partner.

Brains are dazzlingly complex. As a result, there is a growing interest in building artificial neural networks (ANNs) that serve as tractable proxies for understanding the flow of information through real nervous systems. One way to do this is to teach an ANN to perform the same high-level task as an animal. Trained networks can then be studied to give insight into how biological neural circuits work — an approach that has previously been used to help researchers understand how the primate visual system categorizes objects[4] and how the fruit fly detects visual motion[5].

Although this work is exciting, how much trained ANNs can tell scientists about real brains remains unclear. Many different network models can produce the same output, making it difficult to identify one that is better than another. Cowley *et al.* reveal a way to overcome this challenge using the fruit fly *Drosophila melanogaster*. The authors focus on how a male fly chases and sings to a potential mate[6], an activity that relies on visual signals that must pass through a diverse set of lobula columnar (LC) neurons. These neurons form a bottleneck between the retina of the eye and the central brain. LC neurons of particular subtypes converge on small regions of the brain called glomeruli, which contain compact collections of the neuronal processes (axons and dendrites) that transmit signals between neurons.

This glomerular organization suggests that different LC subtypes might be separate channels that are tuned to respond to specific visual features, driving particular actions. Indeed, artificial stimulation of one LC subtype that is responsive to visual looming can cause flies to carry out an escape behaviour[2]. However, the degree to which LC subtypes represent separate sensorimotor channels in more natural settings is not so clear: several LC subtypes have been shown to respond to the same visual features and, in some cases, they make connections with the same downstream neurons[2].

To investigate how LC neurons contribute to courtship decisions, the authors trained ANNs that were optimized for the specific task of accurately predicting a male fly's behaviour when presented with an approximate image of what the male would see during courtship. In a novel twist, they further constrained their model using an approach they call 'knockout training'. Unlike classic machine-learning techniques that 'silence' random artificial units during training to limit overfitting — a problem whereby a model makes accurate predictions for training data but not for other data — knockout training instead silences specific units during training to predict the behaviour of animals in which real, identified neurons are also genetically silenced. This approach is expected to yield direct mapping between artificial units and real neurons.

The authors applied this technique to a large body of data in which one of 23 LC subtypes is genetically silenced in courting male flies. They then designed a task-optimized ANN that is given abstract images of the female fly, as experienced by the courting male, as its input. These images are processed by a 'vision network' (representing the optic lobe of the fly brain) and then a bottleneck layer of 'LC units' (representing each of the 23 LC subtypes). Finally, the outputs of the LC units are passed

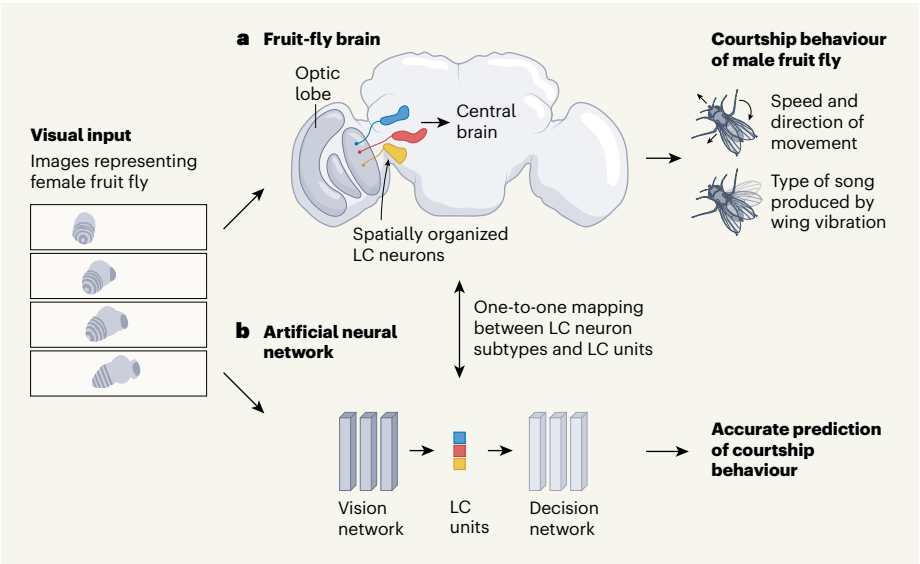

Figure 1 | An artificial neural network that mimics the fruit-fly visual system. **a**, During courtship, a male fruit fly (*Drosophila melanogaster*) exhibits certain behaviours (such as movement and wing vibration) in response to visual information about a female (such as its position and size). In the fly brain, visual information passes from the optic lobe to the central brain through a 'bottleneck' of highly organized lobula columnar (LC) neurons, of which there are 23 distinct subtypes (only 3 subtypes are shown). **b**, Cowley *et al.*[3] trained an artificial neural network to represent the fly visual system: a vision network passes information to a decision network through a bottleneck of 23 LC units, each of which has been trained to correspond directly to one LC subtype. The trained model not only succeeded in predicting the fly's behaviour on the basis of visual input, but also revealed that each LC subtype responds to more than one visual feature and is responsible for more than one behaviour.

through a decision network (representing the central brain) that predicts what the real male fly's behaviour would be in response to those images (Fig. 1). Using knockout training, the authors generated a range of successful networks, each of which showed one-to-one mapping between artificial LC units and known LC neuron subtypes. The authors found that artificial and real fly LC neurons respond similarly to abstract and naturalistic visual stimuli, partially validating the explanatory power of their best networks.

By 'looking under the hood' of successfully trained networks, Cowley *et al.* found that LC units encode visual information about the courted female in a combinatorial manner – that is, with highly overlapping visual tuning. They also observed that LC units regulate male courtship in a distributed and redundant way: many LC units must be silenced to profoundly disrupt behavioural predictions. In support of these network-based findings, a graph of all of the connections between neurons in the real fly brain, referred to as the connectome, shows that LC neuron subtypes share many visual inputs and also fan out to multiple, overlapping central brain regions.

As scientists are often reminded, "all models are wrong but some are useful"[7]. Therefore, perhaps the greatest value of knockout training comes from the predictions it generates for which LC subtypes are expected to drive specific aspects of courtship behaviour, such as whether the male fly's wing vibrations create a song that is constant or pulsing. These predictions should be tested in future laboratory experiments.

To increase the quality of predictions made by the network, knockout training could be improved in several ways. Unlike the feedforward ANNs used in this study, LC neurons (and many brain circuits) are highly recurrent, meaning that they receive feedback from downstream areas. Moreover, sensorimotor mapping can be modulated[8] by an animal's ongoing behaviours[9]. If ANNs were to take neural and behavioural feedback into account, this might greatly improve what they can tell neuroscientists about the brain.

This study shows how precise, large-scale neural-perturbation data can be used to improve the interpretability of artificial models of the brain. The scope of this study might seem narrow at first – the authors focus on only one class of fly visual neurons during a single behavioural task. However, because distributed neural encoding improves robustness, multitasking and efficiency, it is likely to be found across species, including humans. Therefore, something similar to knockout training might ultimately be necessary to understand larger brains, such as those of rodents. How this might be accomplished in animals in which neurons cannot be identified or silenced with such precision remains unclear. For now, as in the past, the simple fruit fly can help by illuminating the way forwards.

Pavan Ramdya is in the Brain Mind Institute and the Institute of Bioengineering, Swiss Federal Institute of Technology in Lausanne (EPFL), Lausanne CH-1015, Switzerland.
e-mail: pavan.ramdya@epfl.ch

1. Coen, P. *et al. Nature* **507**, 233–237 (2014).
2. Wu, M. *et al. eLife* **5**, e21022 (2016).
3. Cowley, B. R. *et al. Nature* **629**, 1100–1108 (2024).
4. Yamins, D. L. K. *et al. Proc. Natl Acad. Sci. USA* **111**, 8619–8924 (2014).
5. Mano, O., Creamer, M. S., Badwan, B. A. & Clark, D. A. *Curr. Biol.* **31**, 4062–4075 (2021).
6. Dickson, B. J. *Science* **322**, 904–909 (2008).
7. Box, G. E. P. in *Robustness in Statistics: Proceedings of a Workshop* (eds Launer, R. L. & Wilkinson, G. N.) 201–236 (Academic, 1979).
8. Turner, M. H., Krieger, A., Pang, M. M. & Clandinin, T. R. *eLife* **11**, e82587 (2022).
9. Chen, C. L. *et al. Nature Neuroscience* **26**, 682–695 (2023).

The author declares no competing interests.
This article was published online on 22 May 2024.

Drug discovery

Obesity drug rewires brain circuits for appetite

Tyler M. Cook & Darleen Sandoval

A two-in-one drug that modulates neural pathways involved in appetite and reward might prove to be more effective and longer lasting than current weight-loss drugs on the market. **See p.1133**

The brain is responsible for regulating how much is eaten and how many calories are burned throughout the day, but in obesity, this balance is disrupted, causing weight gain. Until a few years ago, the most effective strategy for sustained weight loss was surgery. Now, the popular anti-obesity drugs semaglutide (sold under the names Ozempic and Wegovy) and tirzepatide (sold as Mounjaro) are becoming almost as effective as surgery at evoking weight loss. These drugs are based on a peptide made in the gut called GLP-1, and they act on GLP-1 receptors in the brain to suppress eating. However, these medications are costly, in high demand and produce variable results – driving the need for more obesity treatment options. On page 1133, Petersen *et al.*[1] describe how combining a GLP-1 receptor activator with another drug that acts on the brain could be an effective treatment for obesity.

Activators (agonists) of the GLP-1 receptor have been shown to stimulate neurons that express the receptor in areas of the brain that are exposed to the blood circulation – namely the brainstem and the hypothalamus[2]. This leads to the activation of neural circuits in higher brain centres that regulate eating, appetite and reward.

Petersen and colleagues designed a drug that consists of a GLP-1 receptor agonist linked to an inhibitor (antagonist) of another receptor found throughout the brain called the NMDA receptor. This receptor binds to the neurotransmitter molecule glutamate and has an important role in regulating synaptic plasticity – a process that allows communication between neurons to adapt to patterns of activity by strengthening or weakening neuronal connections. NMDA receptors and glutamate signalling have also been linked to obesity in human genome studies, strengthening the rationale for this pharmacological strategy[3].

Although targeting NMDA receptors for weight loss is not a new idea, previous attempts failed because of side effects such as hyperthermia and hyperlocomotion (excessive movement) that occur when NMDA receptors in the brain are inhibited indiscriminately[4]. The authors' dual-mode compound overcomes this problem: the NMDA-receptor antagonist (MK-801, also known as dizocilpine) is activated only after the drug has bound to the GLP-1 receptor and is internalized by the cell. MK-801 can then dampen neuronal excitability and, because it is combined with the GLP-1 receptor agonist, this action is isolated to neurons that express the GLP-1 receptor (Fig. 1).

The authors found that, compared with a GLP-1 analogue alone and other GLP-1 receptor agonists (including semaglutide), GLP-1–MK-801 was similarly processed by the body but was better at reducing body weight in both rats and mice. Petersen *et al.* showed that focusing NMDA-receptor inhibition specifically on neurons that express the GLP-1 receptor was necessary for this effect, whereas linking MK-801 to other gut peptides