chicken coops, greenhouses and so on. Such manual verification can impair reproducibility, and is therefore unusual in conventional applications of machine learning. However, it might be necessary when machine learning is being used to provide data for policymaking. The researchers have made available the data sets that they used to validate and test their method, as well as the resulting identifications. This will be tremendously valuable for further research in this area.

Kruitwagen and colleagues' results provide detailed information about many more solar installations than were previously described. On the basis of their detections, they also estimate the total global electricity-generation capacity of installed facilities, which is very close to the aggregate value provided by the International Energy Agency. Demonstrating the value of their granular facility-level data, they show that non-residential photovoltaic installations were most often developed on agricultural land, indicating a possible trade-off between renewable-energy development and food supply.

So what are the next directions for this field, now that Kruitwagen *et al.* have taken it to the global scale? One focus could be to increase the performance of this approach when identifying installations in different contexts. For example, landscapes — and the photovoltaic installations in them — can look very different from space in different geographical contexts, and so a model trained on data from one region might not make good predictions for another region. Ensuring that the machine-learning system does not systematically underperform for certain regions was a challenge for Kruitwagen and colleagues, and constitutes a problem to be addressed in the future. The same holds for other factors that determine what solar panels and their environments look like, such as types of building, technological characteristics, and the differences between urban and rural environments. This might become even more challenging when the goal is also to detect smaller-scale (residential) solar installations that were not covered in the authors' study.

Another task will be to measure the costs and benefits of the machine-learning approach. Initial concerns that machine learning would require too much effort and would result in low accuracy — and therefore would not add value to conventional approaches for collecting these data — have now been mitigated by initial successes. That said, high-resolution satellite images are expensive, and so is the large amount of computational power needed to train and deploy the model (for example, 71 megawatt hours of energy were used for Kruitwagen and co-workers' study). The costs and benefits to society of such approaches should therefore be evaluated regularly, and compared with the alternative option of collecting data directly.

In this regard, Kruitwagen and colleagues' study is a superb example of the value of machine-learning approaches — researchers, policymakers, international organizations and more across the world need access to trustworthy, regularly updated global data sets such as these. So who should bear the costs of this type of work and ensure that it is maintained and accessible to relevant stakeholders in the long run?

Research laboratories in universities have played a large part in developing similar machine-learning approaches focused on questions of societal relevance, for example to address climate change[10]. Universities and research institutions are ideally suited to conduct such work, but are unlikely to be the right kind of stakeholder to maintain and update large-scale databases over time. Instead, national and international agencies, not-for-profit organizations or publicly funded research entities could have the capacity, sustained funding and public-interest mandate to maintain the infrastructure needed to turn these data into a public good.

Ultimately, analysis of photovoltaic installations at large geographical scales is needed for real-world impact. Kruitwagen *et al.* show that machine learning offers an attractive option for gathering data at such scales when information is not collected through other means.

Lynn H. Kaack is in the Data Science Lab, Hertie School, 10117 Berlin, Germany.
e-mail: kaack@hertie-school.org

1. Kruitwagen, L. *et al. Nature* **598**, 604–610 (2021).
2. Sunter, D. A., Castellanos, S. & Kammen, D. M. *Nature Sustain.* **2**, 71–76 (2019).
3. Barbose, B., Darghouth, N., O'Shaughnessy, E. & Forrester, S. *Tracking the Sun, 2021 Edition* (Lawrence Berkeley Natl Lab, 2021); https://emp.lbl.gov/tracking-the-sun
4. Malof, J. M., Bradbury, K., Collins, L. M. & Newell, R. G. *Appl. Energy* **183**, 229–240 (2016).
5. Yuan, J., Yang, H.-H. L., Omitaomu, O. A. & Bhaduri, B. L. in *2016 Int. Conf. Big Data* 2703–2708 (IEEE, 2016).
6. Ishii, T. *et al.* in *IEEE 23rd Int. Conf. Pattern Recognition* 3344–3349 (IEEE, 2016).
7. Yu, J., Wang, Z., Majumdar, A. & Rajagopal, R. *Joule* **2**, 2605–2617 (2018).
8. Imamoglu, N., Kimura, M., Miyamoto, H., Fujita, A. & Nakamura, R. Preprint at https://arxiv.org/abs/1704.06410 (2017).
9. Hou, X., Wang, B., Hu, W., Yin, L. & Wu, H. Preprint at https://arxiv.org/abs/1912.03685 (2019).
10. Rolnick, D. *et al.* Preprint at https://arxiv.org/abs/1906.05433 (2019).

The author declares no competing interests.

Neuroscience

A brain signal that makes mice hungrier for reward

Lola Welsch & Brigitte L. Kieffer

Release of opioid peptide in the brain leads food-deprived mice to eat more sugar than do mice that are well fed. This opioid signalling mechanism fine-tunes the reward value of food according to the animal's state. **See p.646**

Eating even a simple snack is much more pleasurable when you are hungry than when you are already well fed. During fasting, complex brain mechanisms evaluate an animal's internal state, as well as the caloric and pleasurable (hedonic) values of the food, to ultimately drive eating. In the brain, several molecular signalling systems act together in this process, including the opioid system, which is composed of several opioid peptides and their receptors. The latter are the targets of opioid drugs such as morphine and heroin that have strong pain-reducing and notoriously addictive properties[1,2]. The brain's opioid system contributes to the hedonic value of natural rewards such as food, sex and social interactions[3], but the exact opioid peptide signal and receptor involved, and where they interact, have been challenging to determine. On page 646, Castro *et al.*[4] identify an opioid-system-tuned brain circuit that drives hungry mice to eat more of a sugar reward than do well-fed mice.

The study by Castro and colleagues showed that mice that were deprived of food for 24 hours ate, on average, about three times as many sugar pellets as did mice that had had free access to food. The team found that, in hungry mice, opioid peptides called enkephalins are produced in a part of the brain that is central to reward processing, the nucleus accumbens (NAC)[5], where they bind to and activate local μ-opioid receptors. These receptors are present at the terminals (endings of projections) of incoming neurons that originate from the dorsal raphe nucleus (DRN), which is known as the main centre for mood control. The enkephalin–receptor interaction in the NAC blocks the activity of these incoming neurons, interrupting a communication

path between the two brain structures; as a consequence, mice eat more sucrose pellets. This mechanism (Fig. 1) explains at least half of the hunger-potentiated eating response, and it is almost absent in a satiated mouse.

The precision with which Castro and colleagues characterize this signalling mechanism reflects a technical tour de force for the opioid field and, more broadly, for research into neuropeptides, which are difficult to study. About 20 opioid peptides are produced throughout the brain, and are classified into 3 families: enkephalins, endorphins and dynorphins. These peptides can activate three types of opioid receptor (called μ, δ and κ) that are distributed across various brain circuits, making specific peptide–receptor interactions difficult to tease out. In addition, opioid peptides are rapidly broken down in the body, and the interaction between peptides and their receptors is highly transient, making opioid-peptide function extremely challenging to investigate[6]. Indeed, Castro *et al.* used an impressive armada of genetic tools in mice to characterize the signalling mechanism and its effects on behaviour.

Briefly, experiments in which a drug that blocks opioid receptors was injected into the NAC, and experiments in mice engineered to lack the gene that encodes the μ-opioid receptor, showed that this particular receptor is the key player in enhancing sucrose intake. However, stopping μ-opioid receptor expression specifically in cells of the NAC did not change the effect of food deprivation on sugar consumption, suggesting that the receptor involved is made by neurons projecting from another brain site. Castro *et al.* therefore deleted the receptor gene in three brain regions that contain neurons projecting to the NAC, and were able to identify the DRN as the origin of the receptor. Deletion and re-expression of μ-opioid receptors in the population of DRN neurons that project to the NAC, as well as the rescue of receptor signalling in these neurons, confirmed that μ-opioid receptors expressed at the terminals of DRN–NAC neurons are necessary and sufficient to explain half of the effect of food deprivation on sugar consumption.

Another notable step was the use of an innovative approach, involving the CRISPR gene-editing technique, to block the expression of enkephalin in the NAC. This showed that enkephalins produced in the NAC neurons are necessary for the effects of hunger on feeding. This approach could be developed to identify the roles of other neuropeptides in various brain circuits.

This exemplary high-resolution study, one of the first of its kind in opioid signalling in the brain, will pave the way for future discoveries concerning how opioid signalling subtly regulates homeostatic processes that set hedonic value. Our limited insights into this field contrast with our much broader knowledge of mechanisms that underlie the potent and often devastating effects of opioid drugs, which hijack the opioid system and compromise its delicate balance.

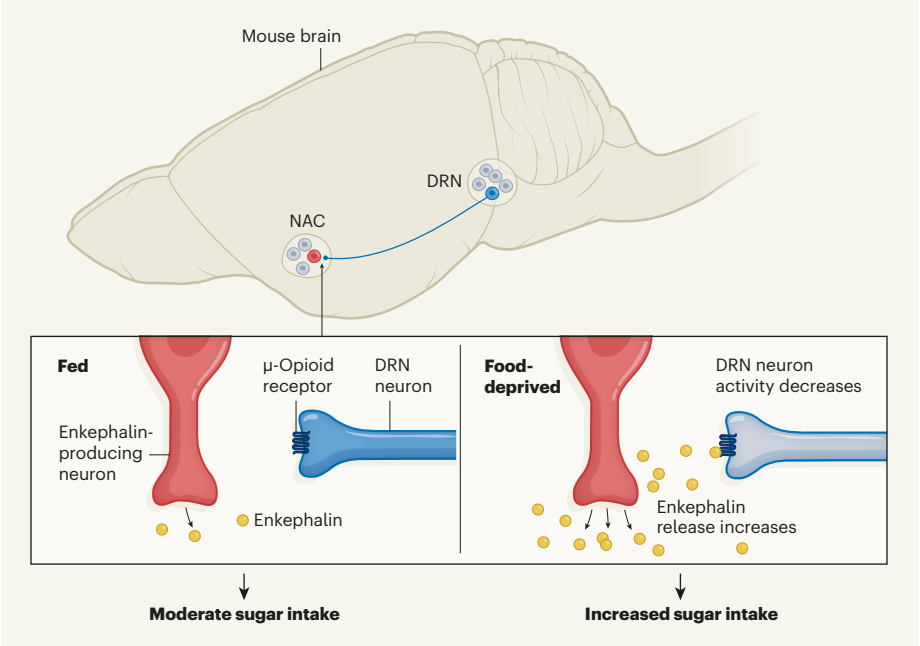

Figure 1 | Opioid signalling in the brain enhances eating in mice deprived of food. Castro *et al.*[4] characterized a brain circuit that allows the adaptation of food intake to the current body state (fed or food-deprived). The circuit involves two brain structures: the dorsal raphe nucleus (DRN), which regulates mood, and the nucleus accumbens (NAC), which is central to reward processing and motivation. Some neurons from the DRN project to the NAC, and the ends of their projections in the NAC express the μ-opioid receptor. With food deprivation, the amount of enkephalin peptides that are released by NAC neurons and that bind to the μ-opioid receptors increases. Activation of the receptors, in turn, inhibits the activity of DRN neurons and leads to food-deprived mice eating more sugar pellets than well-fed mice do.

What have we learnt about reward consumption? The brain integrates internal and external stimuli, and attributes a valence (positive or negative value) to them. The animal subsequently reacts to seek pleasurable experiences and avoid aversive ones. However, the valence of a stimulus is inherently dependent on the body state: in this example, food has a greater positive value when an animal is hungry than when it is well fed. Here, Castro and colleagues demonstrate that enkephalins acting at μ-opioid receptors in the DRN–NAC circuit fine-tune the reward value of sugar depending on the animal's body state (baseline or motivated).

A reasonable assumption is that this mechanism enhances hedonic and motivational values of food[7] when the animal is deprived of food, thereby driving the animal to eat more of the sugar pellets. It is possible that dysregulation of the newly identified interaction between enkephalins and the μ-opioid receptor — caused in humans, for example, by opioid misuse or psychiatric conditions — might produce aberrant adjustment of the reward value of food to the internal body state, jeopardizing the adaptation of feeding behaviour. Thus, these findings probably have implications for eating disorders.

An intriguing aspect of the study is the finding that the μ-opioid receptor-expressing neurons involved in this process originate in the DRN, revealing an opioid mechanism that modulates communication between the NAC and the DRN. The contribution of opioid peptides and receptors to reward processing and motivation is conventionally associated with signalling by the neurotransmitter molecule dopamine, and the NAC is the major brain site in which dopamine is released. Thus, the DRN is a relatively new player in this context.

The DRN produces serotonin, the main neurotransmitter for mood control, and regulates a plethora of processes — from those involved in basic physiology (such as sleep and body temperature) to those influencing food intake, emotional responses and social behaviours. Research in the past few years has also implicated the DRN in reward processing[8], and implies a role for the DRN in the integration of environmental, body and emotional states to gauge the value of rewards[9].

Thus, in the same way as the activity of some DRN neurons promotes sociability in mice in a loneliness-like state[10], the DRN neurons identified by Castro and colleagues enhance eating in animals previously deprived of food. This mechanism bridges hedonic homeostasis and mood in the adaptation of behavioural responses to an internal state. Further studies of how opioid peptides and opioid drugs modulate

the activity of DRN neurons will be crucial, especially given that Castro and colleagues' results indicate that about half of DRN neurons in mice express μ-opioid receptors[4].

Inevitably, questions remain. For example, which mechanisms regulate the other half of the hunger-driven increase in eating? And what are the upstream signals that communicate the hunger state to the NAC? The authors speculate that neurons in a structure called the hypothalamus at the base of the brain, known to produce peptide signals that regulate feeding and project to the NAC, might be involved. If they are right, here comes another microcircuit that plays a feeding-related part in the big orchestra of the brain.

Lola Welsch and **Brigitte L. Kieffer** are in the Department of Psychiatry, McGill University, Montreal, Canada, and at INSERM U1114, University of Strasbourg, Strasbourg 67084, France.
e-mail: brigitte.kieffer@unistra.fr

1. Corder, G., Castro, D. C., Bruchas, M. R. & Scherrer, G. *Annu. Rev. Neurosci.* **41**, 453–473 (2018).
2. Darcq, E. & Kieffer, B. L. *Nature Rev. Neurosci.* **19**, 499–514 (2018).
3. Meier, I. M., Eikemo, M. & Leknes, S. *Curr. Addict. Rep.* **8**, 306–318, (2021).
4. Castro, D. C. *et al. Nature* **598**, 646–651 (2021).
5. Castro, D. C. & Bruchas, M. R. *Neuron* **102**, 529–552 (2019).
6. Al-Hasani, R. *et al. eLife* **7**, e36520 (2018).
7. Castro, D. C. & Berridge, K. C. *J. Neurosci.* **34**, 4239–4250 (2014).
8. Luo, M., Zhou, J. & Liu, Z. *Learn. Mem.* **22**, 452–460 (2015).
9. Luo, M., Li, Y. & Zhong, W. *Neurobiol. Learn. Mem.* **135**, 40–49 (2016).
10. Matthews, G. A. *et al. Cell* **164**, 617–631 (2016).

The authors declare no competing interests.
This article was published online on 13 October 2021.

Immunotherapy

Bacteria recycle tumour waste to fuel immune cells

Laurence C. Chen & Yvonne Y. Chen

Key nutrients that are needed by immune cells are scarce in tumours. Engineered cancer-invading bacteria can recycle tumour waste into metabolic fuel to boost anticancer immune responses in mice. **See p.662**

Harnessing a person's immune system to target a tumour — a type of treatment called cancer immunotherapy — has emerged as a promising treatment option for cancer. T cells are a type of immune cell that has a surveillance function and the capacity to kill foreign or infected cells that are perceived to pose a threat. These properties have cemented T cells as a central pillar of cancer immunotherapy. The generation of an effective antitumour T-cell response depends crucially on the availability of nutrients such as the amino acid L-arginine[1]. However, the tumour microenvironment poses a challenge because it is nutrient-poor[2]. Canale *et al.*[3] show on page 662 that the treatment of mice with metabolically engineered bacteria produces a local, continuous source of L-arginine in the tumour microenvironment that results in strong, long-lasting antitumour T-cell responses when combined with a form of immunotherapy called checkpoint blockade.

It was shown previously[1] that L-arginine supplementation prolongs T-cell survival, enhances the generation of a memory response, and improves tumour-killing efficacy in a mouse model of a skin cancer called melanoma. However, clinical treatment with L-arginine supplementation is not straightforward. Oral administration would require patients to consume impractically large quantities of the amino acid every day, whereas direct injection into a tumour would be possible only for tumours near the surface of the skin, and might be ineffective owing to leakage of the amino acid out of the tumour. Canale and colleagues hypothesized that a strategy to provide a local, continuous supply of L-arginine in the tumour microenvironment would aid T-cell immunotherapy (Fig. 1).

Ammonia is a waste product of metabolism of cancer cells that accumulates in the tumour microenvironment[4], and it can be converted enzymatically into L-arginine. The authors identified two key steps in the L-arginine biosynthesis pathway. These require the protein argR, which suppresses L-arginine biosynthesis, and the enzyme argA. High levels of intracellular L-arginine inhibit the amino acid's production by argA through negative feedback.

Canale and colleagues sought to capitalize on this knowledge when producing genetically engineered microbes. Bacteria can home to tumour microenvironments, colonize these sites and flourish there[5]. The bacterial strain *Escherichia coli* Nissle 1917 (ECN) is harmless and has a long history of medical use in therapeutics, vaccines and diagnostics[6]. The authors genetically manipulated these microbes, deleting the gene that encodes argR and introducing a mutant version of argA that is not inhibited by negative feedback. An ECN strain was thus generated that the authors called L-Arg bacteria, which converted ammonia into L-arginine both *in vitro* and in the tumour microenvironment.

The authors report that injecting L-Arg bacteria into tumour-bearing mice resulted in an increase in tumour-attacking T cells (a type known as conventional T cells) and a decrease in immunosuppressive T cells (a type termed regulatory T cells). This boost in tumour infiltration by conventional T cells complemented the therapeutic effects of checkpoint-blockade immunotherapy. The latter approach used antibodies that target the protein PD-L1, which is found on tumour

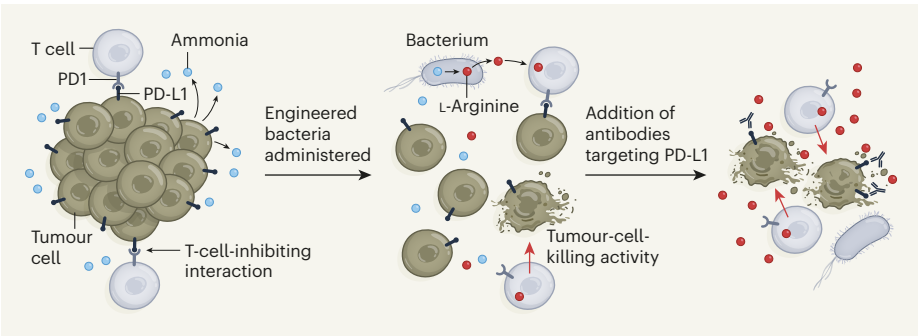

Figure 1 | An approach that boosts antitumour immune responses. Immune cells called T cells face various challenges that hinder their ability to kill tumour cells. The tumour microenvironment is poor in nutrients needed by immune cells, and high levels of metabolic waste products from the tumour cells, such as ammonia, are present. Moreover, T-cell activity can be suppressed as a result of interaction between the PD-1 protein on immune cells and the PD-L1 protein on tumour cells. Canale *et al.*[3] demonstrate that metabolically engineered bacteria that can convert ammonia to the amino acid L-arginine overcome the deficiency of L-arginine in the tumour microenvironment in mice, and thereby boost T-cell function and infiltration into the tumour. The authors report that, when this treatment is combined with a therapy that uses antibodies to block the activity of the PD-L1 protein, the antitumour response is further improved.