researchers have observed activity tracking nonlocal paths (for example, ahead or behind the animal) in the same neurons that usually represent the animal's current location[15]. This may be a direct window into individual trajectories of MB 'mental simulation', and indeed many of the regularities of these nonlocal trajectories are explained by the hypothesis that they are adaptively selected to optimize planning[8].

In the end, perhaps we are not creatures of two minds—or three, or four—but it has become increasingly clear that what we choose depends to a surprising extent on how we compute the values of our candidate actions. And there are many different, interacting routes to this evaluation. ❐

Nathaniel D. Daw
Princeton Neuroscience Institute and Department of Psychology, Princeton University, Princeton, NJ, USA.
e-mail: ndaw@princeton.edu

Published online: 22 October 2018
https://doi.org/10.1038/s41593-018-0258-2

References

1. Daw, N. D., Niv, Y. & Dayan, P. *Nat. Neurosci.* **8**, 1704–1711 (2005).
2. Balleine, B. W. & Dickinson, A. *Neuropharmacology* **37**, 407–419 (1998).
3. Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P. & Dolan, R. J. *Neuron* **69**, 1204–1215 (2011).
4. Lee, S. W., Shimojo, S. & O'Doherty, J. P. *Neuron* **81**, 687–699 (2014).
5. Everitt, B. J. & Robbins, T. W. *Nat. Neurosci.* **8**, 1481–1489 (2005).
6. Gillan, C. M., Kosinski, M., Whelan, R., Phelps, E. A. & Daw, N. D. *eLife* **5**, e11305 (2016).
7. Keramati, M., Dezfouli, A. & Piray, P. *PLoS Comput. Biol.* **7**, e1002055 (2011).
8. Mattar, M. G. & Daw, N. D. *Nat. Neurosci.* https://doi.org/10.1038/s41593-018-0232-z (2018).
9. Dayan, P. *Neural Comput.* **5**, 613–624 (1993).
10. Momennejad, I. et al. *Nat. Hum. Behav.* **1**, 680–692 (2017).
11. Huys, Q. J. et al. *PLoS Comput. Biol.* **8**, e1002410 (2012).
12. Sutton, R. S. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. *Proc. Int. Conf. Mach. Learn.* **7**, 216–224 (1990).
13. Sharpe, M. J. et al. *Nat. Neurosci.* **20**, 735–742 (2017).
14. Miller, K. J., Botvinick, M. M. & Brody, C. D. *Nat. Neurosci.* **20**, 1269–1276 (2017).
15. Foster, D. J. *Annu. Rev. Neurosci.* **40**, 581–602 (2017).

Acknowledgements

I am grateful to my coauthors on the original article, Y. Niv and P. Dayan, and to many other collaborators, trainees, and colleagues who have helped to develop these ideas over the intervening years.

Competing interests

The author declares no competing interests.

NEURAL CIRCUITS

Let it go: central neural control of urination

Two recent studies have expanded our understanding of the circuits controlling urination: one described a projection from brainstem to spinal cord that relaxes the urethral sphincter, and the other revealed a subpopulation of brainstem-projecting layer 5 pyramidal neurons in primary motor cortex that direct the initiation of urination.

Zheyi Ni and Hailan Hu

Although most adults have forgotten the annoyance of incontinence in their first few years, the reality is that disruption of voluntary urination control occurs in at least one out of three people worldwide[1]. Apart from maintaining fluid balance and removing metabolic waste, in other animals urination is also important for social behaviors, like territorial marking and attracting mates. As proposed by Griffith, voiding should meet four criteria: it should be mechanically appropriate (with enough urine in the bladder), safe, emotionally appropriate (regarding embarrassment), and socially appropriate (occurring in the right place and at the right time)[2–4]. Thus, the decision to void or not must be under tight neural control.

Although we have detailed knowledge of the peripheral neural circuits that execute voiding[5–7], much less is known about how it is controlled by the brain. Almost a century ago, Barrington identified a small region in the brainstem, the pontine micturition center (PMC, also known as Barrington's nucleus), as necessary for urination[8]. The PMC transmits signals via the spinal cord to coordinate the contraction of the bladder and relaxation of the urethral sphincter muscles (Fig. 1). Activation of PMC neurons expressing corticotropin-releasing hormone (PMC^{Crh}; ~500 cells in mice) was recently shown to triggers bladder contraction[9].

In a recent study, Keller et al. further identified a small subset (~200 cells) of PMC neurons expressing estrogen receptor 1 (PMC^{Esr1}) as being both necessary and sufficient to induce robust urination by relaxing the urethral sphincter[10]. PMC^{Esr1} and PMC^{Crh} neurons are both excitatory, but have distinct cell body distributions, molecular signatures, and efferent projections. While PMC^{Crh} neurons project more to bladder-contracting preganglionic neurons, PMC^{Esr1} neurons project much more heavily to sphincter-relaxing interneurons in the spinal cord.

To characterize the activity of PMC^{Esr1}neurons in awake, behaving animals, Keller et al. established a quantitative voluntary urination assay in male mice based on their strong motivation to scent-mark in response to female odor. Using the genetically encoded fluorescence Ca^{2+} sensor GCaMP6, the authors performed fiber-photometry recordings and showed that activity of PMC^{Esr1} neurons increased just before the onset of voluntary scent-marking urination and that increased activity continued throughout the voiding process.

Functionally, photostimulation of PMC^{Esr1} neurons expressing channelrhodopsin 2 (ChR2) drove robust urination both in awake and anesthetized animals with high success rate. Recording electrical activities from external urethral sphincter muscle and the bladder pressure showed that while PMC^{Crh} activation solely increased bladder pressure, activation of PMC^{Esr1} neurons induced both bladder contraction and relaxation of urethral sphincter. Furthermore, chemogenetic inhibition of PMC^{Esr1} neurons, but not PMC^{Crh} neurons, abolished voluntary urination. Collectively, these results provide an updated model of how molecularly and functionally distinct cell groups within the PMC control different aspects of voluntary urination (Fig. 1).

While the above studies focused on the subcortical pathways, the involvement of higher brain regions (for example, the cerebral cortex) in urination remains an open question[11]. In this issue of *Nature Neuroscience*, new work shows that a small cluster (~300 cells) of layer 5 (L5) pyramidal neurons in the primary motor cortex (M1) sends commands to the PMC to initiate urination[12]. First, Yao and colleagues injected pseudorabies

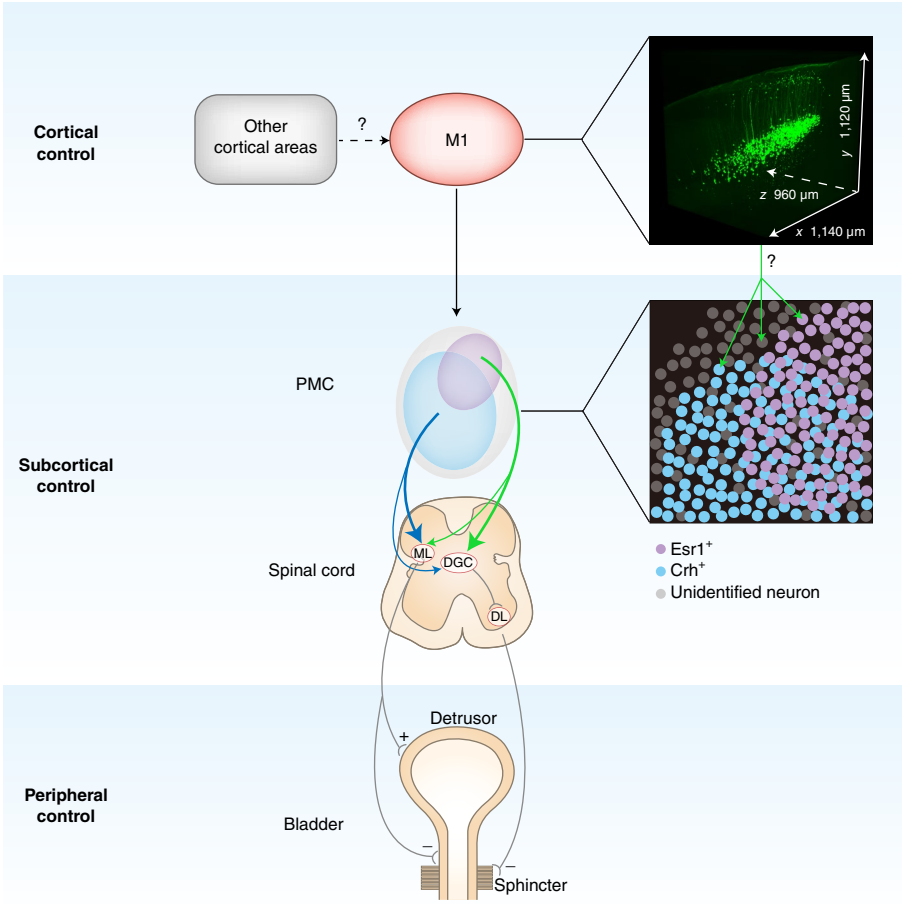

Fig. 1 | Neural circuits for driving micturition. The PMC contains several neuronal types and sends signals to the bladder detrusor or the sphincter via the spinal cord. Recent studies have identified two specific PMC neuronal populations, which express corticotropin-releasing hormone (Crh)[9] or estrogen-1 receptors (Esr1)[10], respectively. Most PMC^{Crh} neurons (blue) project to sacral preganglionic neurons located at the mediolateral column (ML), the activation of which can induce contraction of the bladder detrusor muscle. Most PMC^{Esr1} neurons (pink) project to interneurons in the dorsal gray commissure (DGC) of the spinal cord. Activation of PMC^{Esr1} neurons mainly inhibit the Onuf's nucleus motor neurons in the dorsolateral nucleus (DL), thus driving the urethral sphincter relaxation and urine excretion. The remaining neurons in the PMC that do not express Esr1 or Crh are marked as unidentified neurons (gray), and their roles remain unclear. The PMC receives input from cortical neurons (M1 L5 neurons) to control micturition. Furthermore, the M1–PMC pathway is sufficient to issue the order for micturition. Further experiments are needed to uncover the functions of the other cortical areas during urination. Unconfirmed connections and their functions in micturition are labeled with question marks.

virus 152 encoding EGFP[13] into the bladder walls of adult male mice. This labeled M1 and primary somatosensory cortex. To determine which of these two cortical regions was the direct control unit for driving urination, the authors monitored neuronal population activities in each region in freely behaving mice using fiber photometry[14] and found a robust and transient signal in M1, but not in somatosensory cortex, occurring immediately before the voiding events.

This suggested that this M1 neuronal population could specifically issue the 'order' to initiate voiding. To test this hypothesis, the authors performed optogenetic stimulation specifically in M1 L5 neurons expressing ChR2 and recorded electrical activities from the detrusor muscle of the bladder or the bladder pressure using cystometry. Stimulating these neurons in fact issued an artificial order to initiate voiding. They also observed increases in both electrical activities of the bladder muscle and bladder pressure in response to this cortical stimulation. Importantly, to record the activities of the bladder-projecting M1 neurons more precisely, the authors used the same retrograde labeling technique to express GCaMP2 specifically in these target neurons. They found that activity levels of these neurons were highly correlated with voiding rather than with body movement, in contrast to the bulk-labeled M1 L5 neuronal population signals, which were tightly correlated with body movement.

The next question was which downstream brain area receives the order and via which pathway. The authors traced the axonal projections of these M1 L5 neurons and discovered dense fibers in the PMC. They then specifically expressed GCaMP6s in the recipient PMC neuron using AAV1-hSyn-Cre/AAV1-CMV-Cre, an effective anterograde trans-synaptic virus[15]. Strikingly, 100% of the micturition events were correlated with Ca^{2+} transients, as measured by fiber photometry in these targeted PMC neurons. Furthermore, by expressing ChR2 specifically in these PMC neurons, which are postsynaptic to M1 L5 neurons, the authors could drive voiding in freely behaving mice by optogenetic stimulation. The M1 L5→PMC circuit was therefore sufficient to issue the order for urination.

The last question was whether the 'high command' unit was required for proper urination control. The authors performed three independent loss-of-function experiments to either inhibit or ablate M1 L5 neurons and subsequently test whether micturition was impaired. They found that micturition behavior was suppressed, as reflected by a reduced frequency of micturition events, with increased latency to the first micturition event and a decreased volume of urine deposited during the observation period. Thus, without the order from the high command M1 L5 neurons, urination was greatly disrupted. Taken together, these findings lend strong support to the idea that M1 sits in the high layers of the micturition-controlling hierarchical network, mediating top-down regulation of micturition (Fig. 1).

These two pieces of work provide exciting potential therapeutic strategies for treating urination-related diseases including urine retention and incontinence. They also raise many new and fascinating questions. What are the other higher cortical areas that input into these M1 L5 neurons? Which type (Crh^+ or $Esr1^+$) of PMC neurons do M1 L5 neurons innervate? How does a transient signal (~0.5 s) in the M1 at urination onset translate to more lasting activity in the PMC (~5 s) throughout the voiding process[10,12]? What mediates this hysteresis effect in the M1–PMC circuit? Finally, how do social

contexts regulate the M1–PMC circuit to control the desire for urination? While Hou et al. found that one input of the PMC (i.e., the medial preoptic area) modulates social-hierarchy-dependent micturition patterns[9], here neuronal activities of both M1 L5 neurons and PMC^{Esr1} are shown to increase strongly as male mice intentionally mark their territory[10,12]. It will be fascinating to understand how these different circuits act in synergy to regulate voluntary micturition induced by social signals. ❐

Zheyi Ni and Hailan Hu*
Department of Psychiatry of First Affiliated Hospital, Center for Neuroscience, NHC and CAMS Key Laboratory of Medical Neurobiology, School of Medicine, Zhejiang University, Hangzhou, China.
**e-mail: huhailan@zju.edu.cn*

Published online: 25 October 2018
https://doi.org/10.1038/s41593-018-0259-1

References

1. Minassian, V. A., Drutz, H. P. & Al-Badr, A. *Int. J. Gynaecol. Obstet.* **82**, 327–338 (2003).
2. Fowler, C. J., Griffiths, D. & de Groat, W. C. *Nat. Rev. Neurosci.* **9**, 453–466 (2008).
3. Griffiths, D. *Nat. Rev. Urol.* **12**, 695–705 (2015).
4. Griffiths, D. J. *Curr. Urol. Rep.* **5**, 348–352 (2004).
5. de Groat, W. C., Griffiths, D. & Yoshimura, N. *Compr. Physiol.* **5**, 327–396 (2015).
6. de Groat, W. C. & Wickens, C. *Acta Physiol. (Oxf.)* **207**, 66–84 (2013).
7. Merrill, L., Gonzalez, E. J., Girard, B. M. & Vizzard, M. A. *Nat. Rev. Urol.* **13**, 193–204 (2016).
8. Barrington, F. J. F. *Q. J. Exp. Physiol.* **15**, 81–102 (1925).
9. Hou, X. H. et al. *Cell* **167**, 73–86 (2016).
10. Keller, J. A. et al. *Nat. Neurosci.* **21**, 1229–1238 (2018).
11. Malykhina, A. P. *eLife* **6**, e33219 (2017).
12. Yao, J. et al. *Nat. Neurosci.* https://doi.org/10.1038/s41593-018-256-4 (2018).
13. Smith, B. N. et al. *Proc. Natl. Acad. Sci. USA* **97**, 9264–9269 (2000).
14. Adelsberger, H., Garaschuk, O. & Konnerth, A. *Nat. Neurosci.* **8**, 988–990 (2005).
15. Zingg, B. et al. *Neuron* **93**, 33–47 (2017).

Competing interests

The authors declare no competing interests.

HIPPOCAMPAL REPLAY

Spoiled for choice, pressed for time

A new theory derives the sequential nature of hippocampal replay from first principles and, moreover, predicts the specific patterns of replay that are actually observed in multiple different experiments.

John Widloski and David J. Foster

Hippocampal place cells, the subject of the 2014 Nobel Prize in Physiology, live a secret life. Originally thought to fire action potentials dutifully only within their place fields, they revel in periods of promiscuous propagation. When a rat pauses during exploration of a maze, its place cells, initially paused in their labor as faithful reporters of the animal's current location, suddenly come alive in bursts of activity, which zip sequentially from cell to cell up to 20 times faster than normal. One moment, a sequence depicts a series of places the animal is about to visit, as if rehearsing the journey. The next moment, a sequence travels just as swiftly backwards through the animal's past, as if ruminating over the choices it has made. These activity patterns in the rat's brain are commonly referred to as awake replay, and in this issue of *Nature Neuroscience*, Mattar and Daw present the first theoretical account of why replay patterns take the forms that they do[1].

For behavioral neurophysiologists, stumbling onto the phenomenology of awake replay has been like falling into a box of chocolates. After first discovering that awake replay goes backwards[2], and then that it can go forward as well as backwards[3], it was discovered that it can go both ways at a choice-point in a maze[4,5], and it can also join together different experiences to find shortcuts[4]. Awake replay contributes to decisions[6,7] and can depict the precise trajectory that the animal is about to take all the way to a remembered goal location[8]. It is also exquisitely sensitive to the learned shape of the maze the rat is running on[9], and when a rat discovers unexpected changes in reward, there are corresponding changes in the numbers of awake replays that get produced[10,11].

Can you have too much of a good thing? This assortment of results has exposed the absence of a theoretical framework to make sense of all the data. For example, the distinction between forward and backwards replay has been confusing, with some researchers ascribing them different roles and others preferring to ignore backwards replay altogether. Now, in work of extraordinary elegance, Mattar and Daw provide exactly the sort of theoretical framework that the field has been looking for. There are two major accomplishments. First, they derive replay from first principles, giving what is sometimes called a 'normative' account. That is, they start from the Darwinian injunction–eat and don't get eaten–and from there derive replay sequences as the optimal order in which to sample and learn from place memories to maximize future rewards and minimize future costs. Second, they demonstrate that their framework can account for almost all of the results discovered about replay in the last decade. Taken altogether, it is an astonishing achievement.

So how do they do it? They begin by defining the fundamental unit of experience as a movement between two neighboring locations, given an action choice at the first location and with a resulting outcome. Animals use such experiences to learn to improve the action choices they make, to increase the amounts of reward they will obtain in an environment. By using an algorithm called Q learning[12], they model the outcomes not just as immediate rewards or costs, but also including the long-term expectations about what returns will accrue in the future. Every time a unit of experience is used by the Q learning algorithm to make an incremental improvement in action choices, this is called a 'backup'. During behavior, the backups can be made from the actual moves the animal makes through the world. But backups can also be made offline, that is, when the animal is not actively experiencing the movements but rather replaying them while being paused somewhere else. Mattar and Daw are agnostic about whether offline backups are specific previous experiences recalled from memory or simulated experiences using an internal model of the world. The key question they ask is: given the short period of time that an animal pauses in a maze and the very large number of possible