

NEUROSCIENCE

Volitional activation of remote place representations with a hippocampal brain–machine interface

Chongxi Lai^{1*†‡}, Shinsuke Tanaka^{1†‡}, Timothy D. Harris¹, Albert K. Lee^{1,2*}

The hippocampus is critical for recollecting and imagining experiences. This is believed to involve voluntarily drawing from hippocampal memory representations of people, events, and places, including maplike representations of familiar environments. However, whether representations in such “cognitive maps” can be volitionally accessed is unknown. We developed a brain–machine interface to test whether rats can do so by controlling their hippocampal activity in a flexible, goal-directed, and model-based manner. We found that rats can efficiently navigate or direct objects to arbitrary goal locations within a virtual reality arena solely by activating and sustaining appropriate hippocampal representations of remote places. This provides insight into the mechanisms underlying episodic memory recall, mental simulation and planning, and imagination and opens up possibilities for high-level neural prosthetics that use hippocampal representations.

The ability to simulate scenarios in one’s mind is a hallmark of intelligence, as it facilitates the evaluation of past experiences and future plans. For instance, we can imagine walking around our previous workplace, or imagine how our current workplace might function if we rearranged the furniture. Such imagination requires an internal world model that can be flexibly accessed to construct possible scenarios (1–3).

The hippocampus is a brain region that is critical for memory and imagination (1, 4–6). It holds a model of the environment (also called a cognitive map) (7, 8) that could potentially be mentally traversed for the purpose of recall or simulation. In particular, the hippocampus contains spatial maplike representations of previously explored environments. Each environment’s representation consists of place cells—neurons that fire selectively whenever an animal moves through specific locations (called the “place fields” of those cells) in that environment (9, 10). This selective firing results in a distinct multicell activity pattern at each location in the environment, which, during physical navigation, can be used to decode the animal’s current location from the ongoing pattern of neural activity (11). In contrast, a key aspect of imagination is the activation of neural representations that deviate from current sensory input, such as those that are nonlocal (i.e., represent locations away from one’s current location). Previous work has shown brief and intermittent activation of nonlocal

hippocampal spatial representations suggestive of the planning of specific paths within a cognitive map (12–21). However, it is unknown whether this activity is volitionally controlled or rather reflects passive memory-related processes that are presumably nonvolitional (22, 23).

To test whether an animal can directly control its hippocampal activity according to its model of the world, we used a brain–machine interface (BMI) approach because, unlike with humans, we cannot simply ask animals to imagine scenarios. With BMI methods, we could reward animals for generating neural activity resembling the simulation of specific scenarios. More precisely, we could reward them for the volitional activation of specific nonlocal representations from the cognitive map—a fundamental building block of scenario simulation. BMI research has a rich history of directly testing for volitional control of activity patterns of neuronal ensembles in the motor cortex and related areas (24–35). In the hippocampus, it has been shown that the activity level of individual neurons (36, 37) or the population activity related to individual stimuli (38) can be controlled. However, a real-time BMI that allows humans or animals to control their hippocampal population activity in terms of the content of their cognitive map (e.g., location representations) has never been demonstrated.

A hippocampal map–based BMI

We designed a real-time hippocampal BMI and two BMI tasks to investigate whether rats could navigate to goals (“jumper” navigation task), or move external objects to goals while remaining stationary (“Jedi” object location control task), within an immersive virtual reality (VR) environment solely by controlling the activity of a population of place cells. Each jumper or Jedi BMI experiment consisted of three phases (Fig. 1A). In phase 1, rats ran to a

succession of arbitrary locations marked tall, visible goal cue placed in a familiar dimensional virtual arena (“running” task). Upon reaching each cue, liquid reward was delivered, the trial ended, and the cue moved to another location for the next trial. Animals were secured in a harness and could freely rotate their body and head direction on top of a spherical treadmill (39) while hippocampal CA1 neural activity was recorded (Fig. 1B, fig. S1, and movie S1). We applied a recently developed field-programmable gate array (FPGA)-based neural signal processor to perform low-latency (1 ms) assignment of extracellular spikes (recorded from 128 channels) to a population of hippocampal units (40, 41). In the running task, treadmill movement updated the animal’s location in the virtual environment, and many hippocampal units (i.e., place units) displayed spatially modulated activity (39, 42–44) (Fig. 1B, blue arrows) similar to that in real-world environments (8–11). In phase 2, the binned spike counts from the most recent 1.5 or 5 s of activity of these place units and the animal trajectory from the running task were used to train a decoder (Fig. 1B, green arrows) that estimates the animal’s current location from the neural data every 100 ms. We used a deep neural network for decoding (fig. S2), allowing the use of data augmentation for training—a method that improves both the decoder’s performance given limited data and its noise robustness. In phase 3, the treadmill was disconnected from the VR system, and the animal’s ability to control its own or an object’s translational movement was limited to controlling its hippocampal activity, which was converted by the decoder into a specific location output every 100 ms (Fig. 1C). Note that the decoder was trained to estimate the animal’s current location in the running task only, not its location in the subsequent BMI tasks, but, during BMI periods, the animal needed to generate activity corresponding to locations away from its current location.

BMI navigation task

In the jumper task, we tested whether animals could navigate to arbitrary goal locations as in the running task, except here by means of BMI-based first-person teleportation. After rats performed the running task for ~40 min (~120 trials) (Fig. 2, A and B, and movie S1), the data were used to train the decoder, which accurately estimated the rat’s current location in the running task [validation set coefficient of determination (R^2) = 0.78 to 0.88] (Fig. 2C). Jumper trials were identical to running trials, except the animal’s location was updated to the BMI-decoded location (smoothed with a 3-s sliding window to help reduce potential high-frequency visual jitter of the VR updates) (Fig. 2, D and E, and movie S1). If an animal did not reach the goal within 62 s, the trial

¹Janelia Research Campus, Howard Hughes Medical Institute, Ashburn, VA, USA. ²Howard Hughes Medical Institute and Department of Neurology, Beth Israel Deaconess Medical Center, Boston, MA, USA.

*Corresponding author. Email: chongxi.lai@gmail.com (C.L.); aleec31@bidmc.harvard.edu (A.K.L.)

†Present address: Howard Hughes Medical Institute and Department of Neurology, Beth Israel Deaconess Medical Center, Boston, MA, USA.

‡These authors contributed equally to this work.



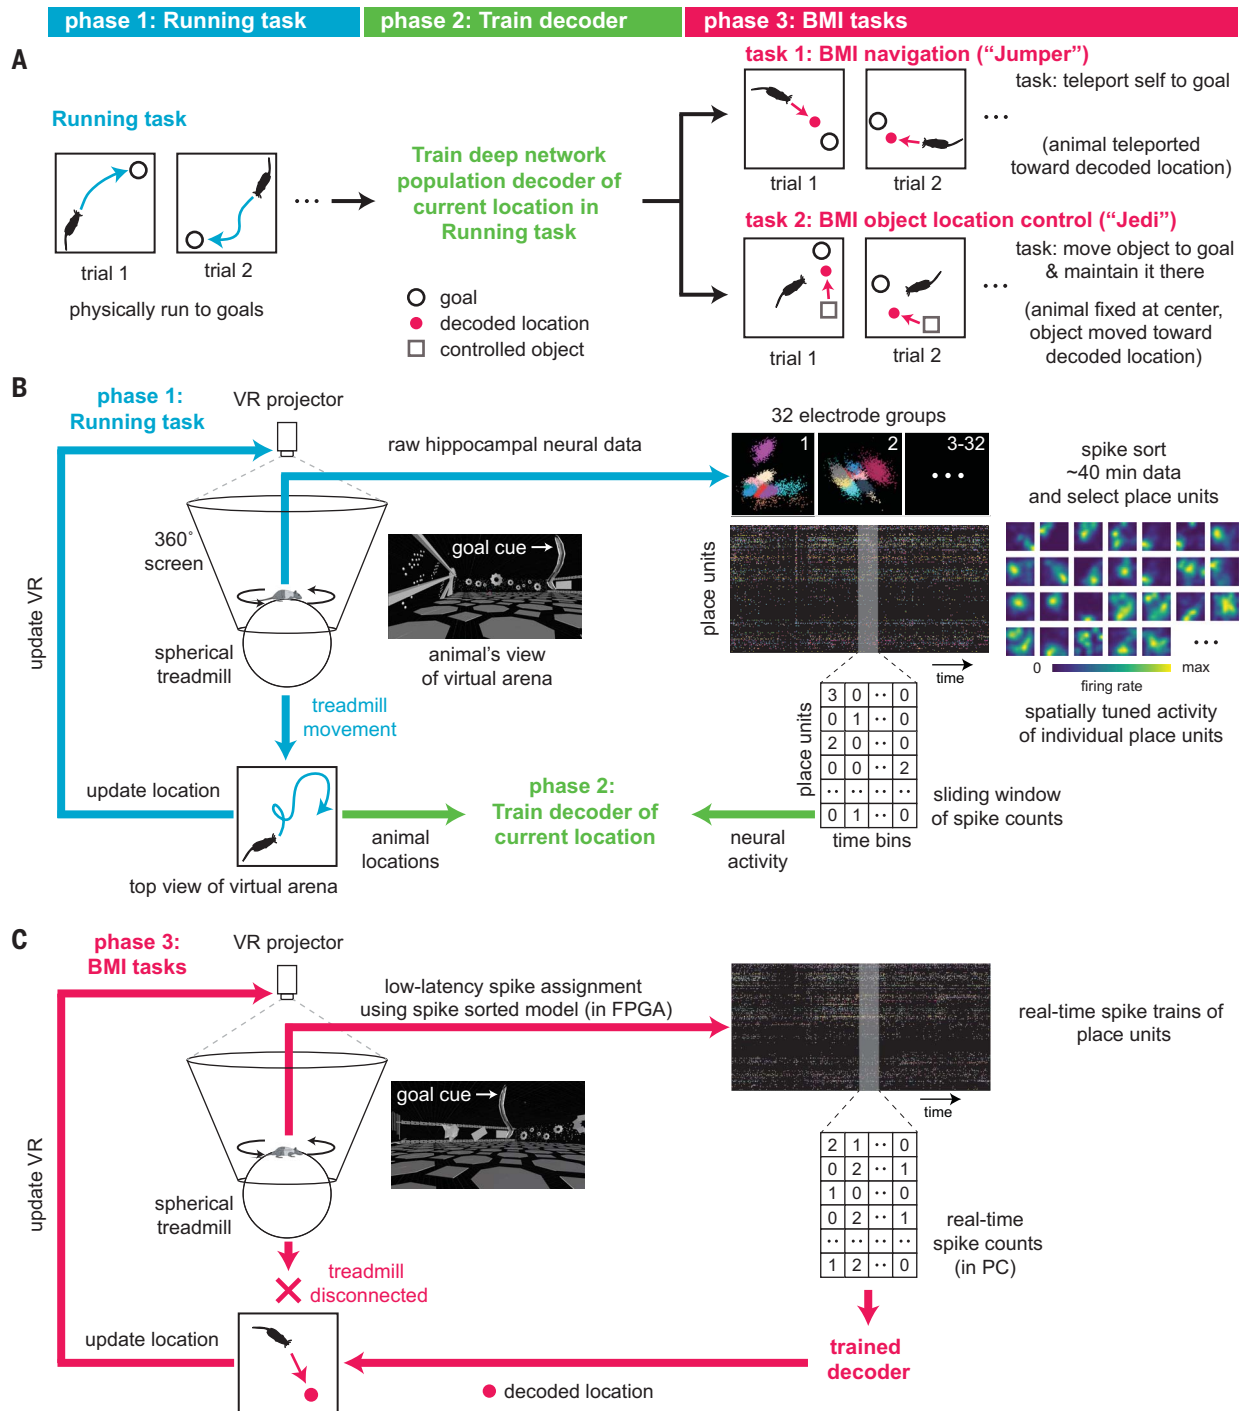


Fig. 1. Hippocampal map-based brain-machine interface in a virtual reality system. (A) Steps for performing the two different BMI experiments in this study. Rats first physically ran to a series of goals (running task), while their hippocampal neural activity and (virtual) location in a square arena were recorded. This data was used to train a decoder to take neural activity as input and output the animal's current location in the running task. In BMI task 1 (jumper task), animals needed to generate neural activity that would be decoded as locations they wanted to move to so that they could reach each goal (to obtain a reward). In BMI task 2 (Jedi task), animals were fixed at the center of the virtual arena (but could rotate) and needed to generate activity corresponding to locations where they wanted an external object to move to so that the object reached the goal, then they needed to sustain that activity to maintain the object there (to maximize reward).

(B) Schematic of the VR system (left). The animal was free to rotate its body in the horizontal plane. In the running task, the animal's location in the virtual arena environment was updated according to treadmill movement. Simultaneously recorded spiking from a population of hippocampal CA1 units expressed place fields—the basis of the cognitive map of the environment (right). Decoder was then trained using binned spiking activity and location data. (C) In both BMI tasks, the treadmill no longer updated VR. Instead, the animal or object location was controlled solely by real-time hippocampal activity. A neural signal processor rapidly assigned activity to individual units, whose spike counts were fed into the decoder. VR projection was updated according to locations output by the decoder. In the jumper (Jedi) task, the animal's (object's) virtual location was moved toward the most recent decoded locations. PC, personal computer.




Fig. 2. Rats can navigate to goals by controlling their hippocampal activity.

In both the running and jumper BMI tasks, animals were rewarded when they reached each goal. **(A)** Animal trajectories in the virtual arena for consecutive running task trials. Trial duration (time to reach goal) in seconds is shown. **(B)** Example running task trial. From top: trajectory, firing rate (z-scored) of individual units (units were ordered by time of peak activity), treadmill speed, and LFP from one recording channel and corresponding wavelet spectrogram during trial. **(C)** Accuracy of trained decoder of animal's current location for held-out running task data. Actual and decoded trajectories during example trial (top left) and across several trials (for x and y coordinates separately, bottom left). Median decoding error (distance between actual and decoded locations) with range and quartiles (bottom right). **(D)** Example jumper BMI trial with similar trajectory as the running trial in **(B)**. From top: trajectory generated

by the animal controlling its hippocampal activity and the decoder output (animal is teleported toward decoded location; each gray circle represents the decoded location at the time the animal is at the corresponding point in the trajectory connected by the dark line, sampled here every 1 s), firing rate of individual units [using same order of units as in **(B)**], treadmill speed, LFP, and spectrogram. **(E)** Example jumper BMI trial in which animal did not move the treadmill. Trajectory (left) as in **(D)**. Right, from top: unit activity, treadmill speed, LFP, and spectrogram. See fig. S10 for all 10 nonmovement trials. **(F)** BMI-generated trajectories for consecutive jumper trials. **(G)** Mean jumper trial duration (magenta vertical line) is significantly lower than distribution of expected mean duration for simulated trials if goals were in random locations. **(H)** Polar distribution of angle between direction of movement and direction to goal during running and jumper tasks. Zero corresponds to animal movement directly toward the goal center.

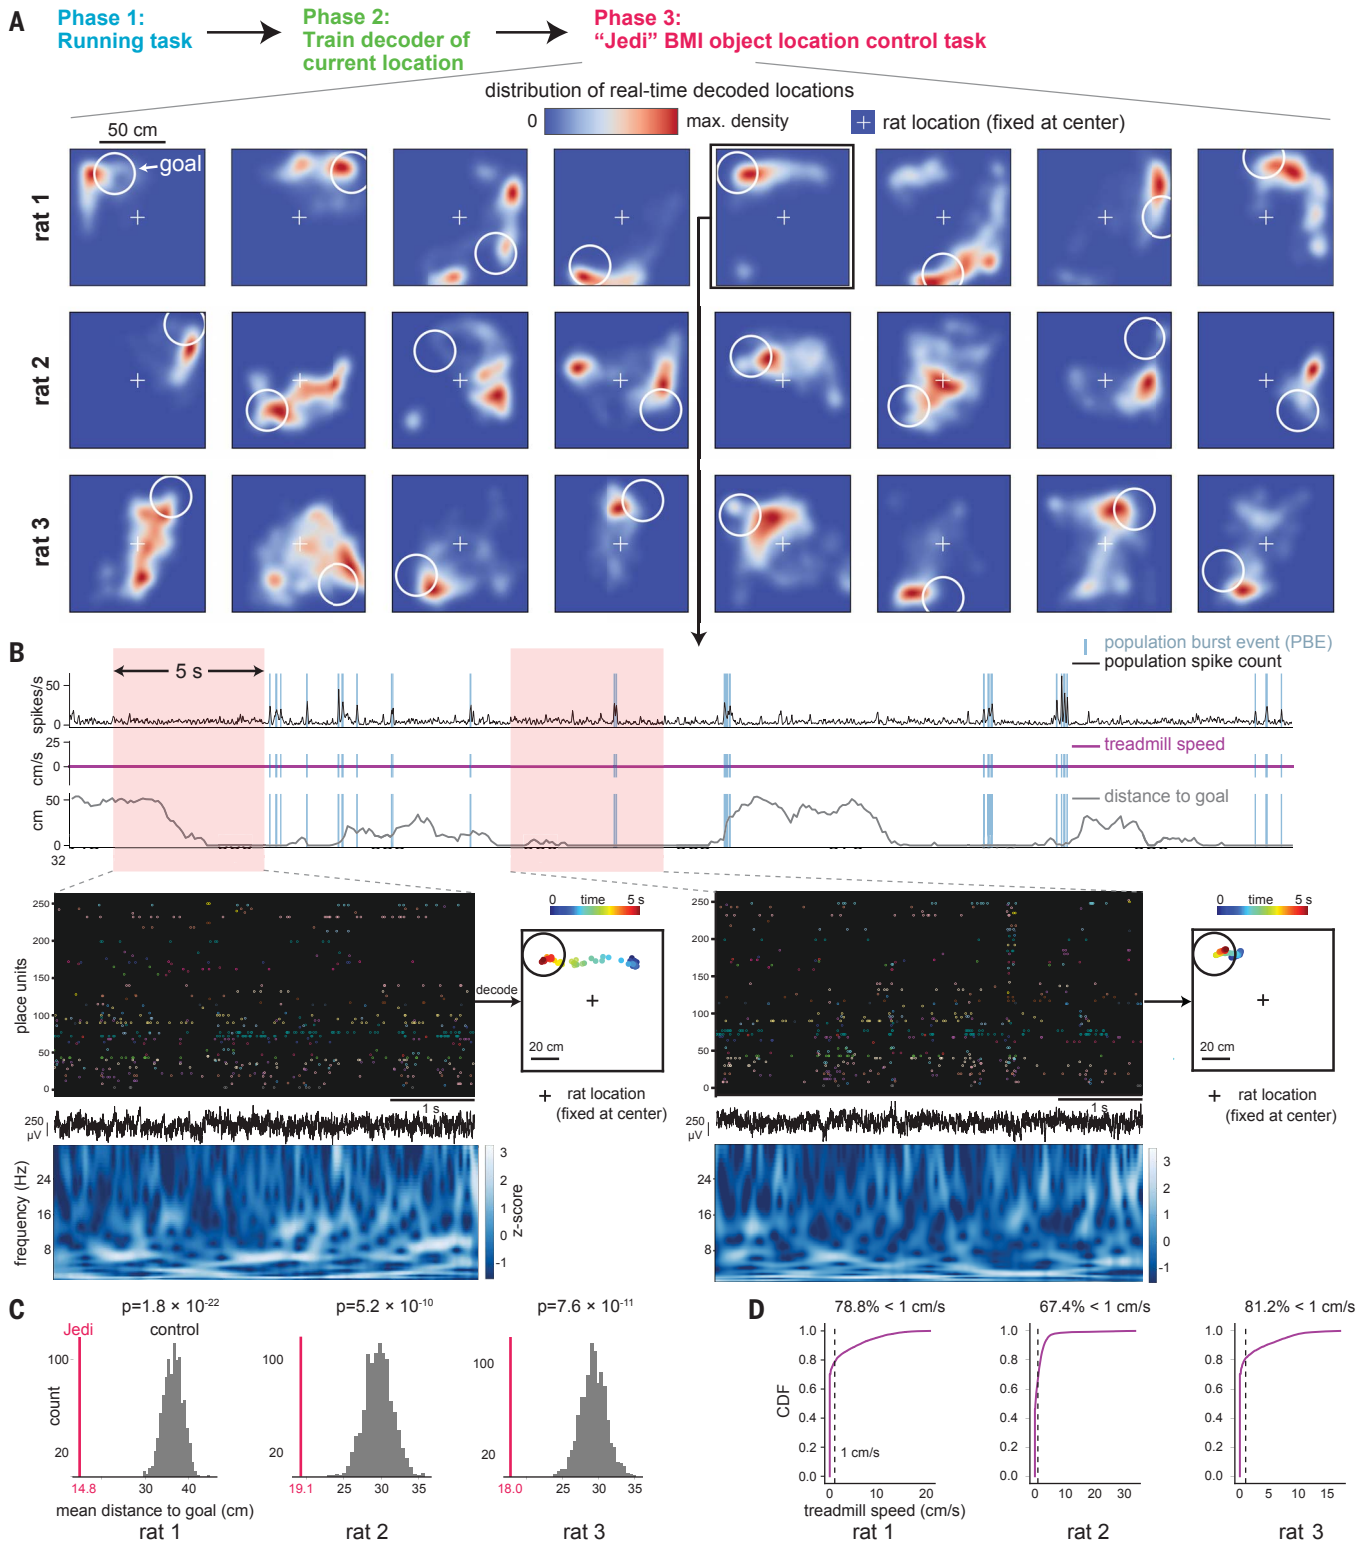


Fig. 3. Rats can move objects to remote goal locations and maintain them there by controlling their hippocampal activity. In the Jedi BMI task, trials did not end when the external controlled object first reached the goal; instead, animals were rewarded as long as the object was in the goal region (white circle), for up to 3 min per trial. The animal was always fixed at center of the virtual arena but could rotate its body and generally turned toward each goal. **(A)** Distribution of real-time decoded locations (output every 100 ms)

generated by the animal controlling its hippocampal activity across eight consecutive Jedi BMI trials for rats 1, 2, and 3. Panels show decoded locations during each trial (up to 3 min; fig. S11). Periods when the animal's body rotated $>12^\circ/\text{s}$ were excluded. See text and methods for details. The external controlled object (which was visible for rats 1 and 2 but invisible for rat 3) was moved toward the decoded location (fig. S11 shows that the distribution of object locations was essentially the same as the distribution of decoded

locations). **(B)** A 40-s-long period during an example trial during which the animal did not move the treadmill. From top: summed activity across all units with PBEs identified, treadmill speed, distance of decoded location from goal (0 means inside goal region), and close-ups of two 5-s periods [as animal moves object to goal (left) and as animal maintains object at goal (right); points in the arena represent sequence of decoded locations] with

spike trains of units, LFP, and spectrogram. See fig. S12 for additional example periods. **(C)** Mean distance of decoded location from goal across all trials (magenta vertical line) is significantly lower than mean distance expected for randomized goal locations. **(D)** Treadmill speed distribution during periods shown in (A), illustrating that the animal was generally still during task performance.

ended and a new goal cue appeared at a random location.

Rats successfully navigated by controlling their hippocampus, generating efficient paths to each goal (Fig. 2F; see figs. S3 to S5 for all trials of three rats, and figs. S6 to S8 for all trials re-decoded using a shorter decoding window and without smoothing). To check whether this performance could be attributed to non-spatially-specific neural activity (e.g., modulating global firing rate), we randomly shuffled the spike trains across place units, ran the shuffled data through the original decoder to produce simulated trajectories, then determined how long it would have taken to reach the same sequence of goal locations as in the original experiment. Shuffled-unit mean trial durations were much longer than the actual means ($P < 10^{-100}$, three rats, one session each), suggesting that performance depended on generating place field-related activity. To test whether generating non-goal-directed sequences of location-specific activity (e.g., random movement within the cognitive map) could explain the performance, we randomly shuffled the goal locations in each trial while preserving the original BMI trajectories and then determined the time that would have been needed to reach the shuffled goals. Shuffled-goal mean trial durations were again much longer than actual means ($P = 2.8 \times 10^{-15}$ to 1.5×10^{-7}) (Fig. 2G), indicating that animals' BMI trajectories were clearly goal-directed. Goal-directedness was also apparent from the distribution of angles between the animal's instantaneous direction of BMI-generated movement and the direction from the animal's current location to the goal, which was concentrated around a value near 0° (Fig. 2H). Thus, even though jumper trials took longer than running trials (mean trial duration across animals: 15.1 s versus 6.9 s; note, however, that BMI decoding and smoothing added a few seconds to jumper durations), the animals' routes revealed effective, goal-directed, map-based BMI navigation. Furthermore, such performance was achieved without extensive BMI training (Fig. 2F; figs. S3 to S5 show sessions 3, 9, and 2 for rats 1, 2, and 3, respectively; table S1; a fourth rat failed to perform either BMI task).

Although animals were free to physically run during the jumper task, such movement was not necessary for task performance. Initially, animals ran as in the running task, but in later trials, animals ran less (fig. S9). In a subset of trials (10 out of 161 trials) (Fig. 2E and fig. S10),

animals remained still, yet in all cases they efficiently reached the goal. Moreover, this successful navigation did not depend on activity in population burst events (PBEs), which often appear during immobility and during which brief activation of place cell representations for remote locations has been shown to occur (13–16, 18, 20, 21, 23).

BMI object location control task

Although episodic memories are encoded and often retrieved using a first-person perspective, individuals can also imagine scenarios from a third-person perspective, with other animate and inanimate players taking part. Furthermore, imagination often involves holding a single thought in mind for extended periods. Therefore, our second BMI task, the Jedi task, tested whether animals—while remaining in the same place—could use the same map of the arena to control the location of a virtual object, guide it to the goal cue location, and maintain it nearby. The jumper and Jedi tasks thus used different forms of feedback: self-location and the location of an object, respectively. After the same running task and decoder training phases as in the jumper experiment, the animals in Jedi were fixed (but could freely turn) at the arena's center, and the object's location was updated to the BMI-decoded location (with a 2-s smoothing window). In each trial, the goal cue remained in the same place, providing reward as long as the object touched it. After 3 min or the rat having received 0.5 ml of reward in total, whichever came first, a new goal cue appeared at a distant random location for the next trial.

Rats could activate and sustain a remote location's representation around the goal for long periods, until the trial ended, and then shift attention to the next goal (Fig. 3, A and B; fig. S11; and movie S1). Performance was measured using the mean distance (over time) between the decoded locations and goals. Shuffling spike trains across units yielded much greater mean distances than the actual means ($P = 2.2 \times 10^{-5}$ to 2.6×10^{-3} , three rats, one session each). To assess the goal-directedness of BMI-generated activity, we shuffled the goal locations while preserving the locations output by the decoder. The decoded (and controlled object's) location was far more concentrated around the actual remote goal cue than shuffled goal locations ($P = 1.8 \times 10^{-22}$ to 5.2×10^{-10}) (Fig. 3C and fig. S11), indicating clear goal-directed control of activity. Again, such per-

formance occurred without extensive training (Fig. 3A shows sessions 7, 6, and 3 for animals 1, 2, and 3, respectively). Task performance was not dependent on PBEs, as there was no change in performance when all activity in PBEs was eliminated and the decoder was rerun post hoc (fig. S11).

Animal movement was generally low when engaged in the Jedi task (Fig. 3D), and movement was not required for successful performance. There were many longer periods (≥ 8 s long with a treadmill speed of ≤ 1 cm/s, 38 periods, mean: 17.3 s, maximum: 44.0 s) during which the animal did not move the treadmill while it directed the object to the goal and/or held it there (34 of 38 periods) (Fig. 3B and fig. S12). Activity during PBEs was also generally not necessary for performance in these non-movement segments (fig. S12).

Features of volitionally generated spiking and local field potential activity

What characteristics did the volitionally generated activity have? First, mean firing rates per unit were similar between jumper and running tasks (fig. S13A). Mean firing rates per unit were correlated across Jedi and running tasks, but lower in Jedi (fig. S13B)—consistent with the decreased physical movement in Jedi.

We then investigated the hypothesis that, to move themselves or the object toward a given (decoded) location in the jumper and Jedi tasks, animals generated a pattern of firing rates across units (i.e., a population vector, or PV) similar to the mean PV at that location over the entire running task (called the reference PV, or rPV) (Fig. 4A). (Note that the set of rPVs for all locations is thus equivalent to the standard place field map across the population.) We examined the correlation between the PV generated at each moment (in every 500-ms window) during jumper or Jedi and the rPV of the decoded location at that moment. As a benchmark, we computed the correlation between the 500-ms PVs during the running task with the rPVs corresponding to the animal's actual location at those times (Fig. 4, B and C, "Run") as well as the correlation between the running task PVs and the rPVs of random locations (Fig. 4, B and C, "randRun"). We then correlated jumper or Jedi PVs with the rPVs of the decoded locations at each moment (Fig. 4B for "jumper" and Fig. 4C and fig. S14 for "Jedi") and with rPVs of random locations ("randJumper" and "randJedi"). Jumper and Jedi PVs were significantly correlated

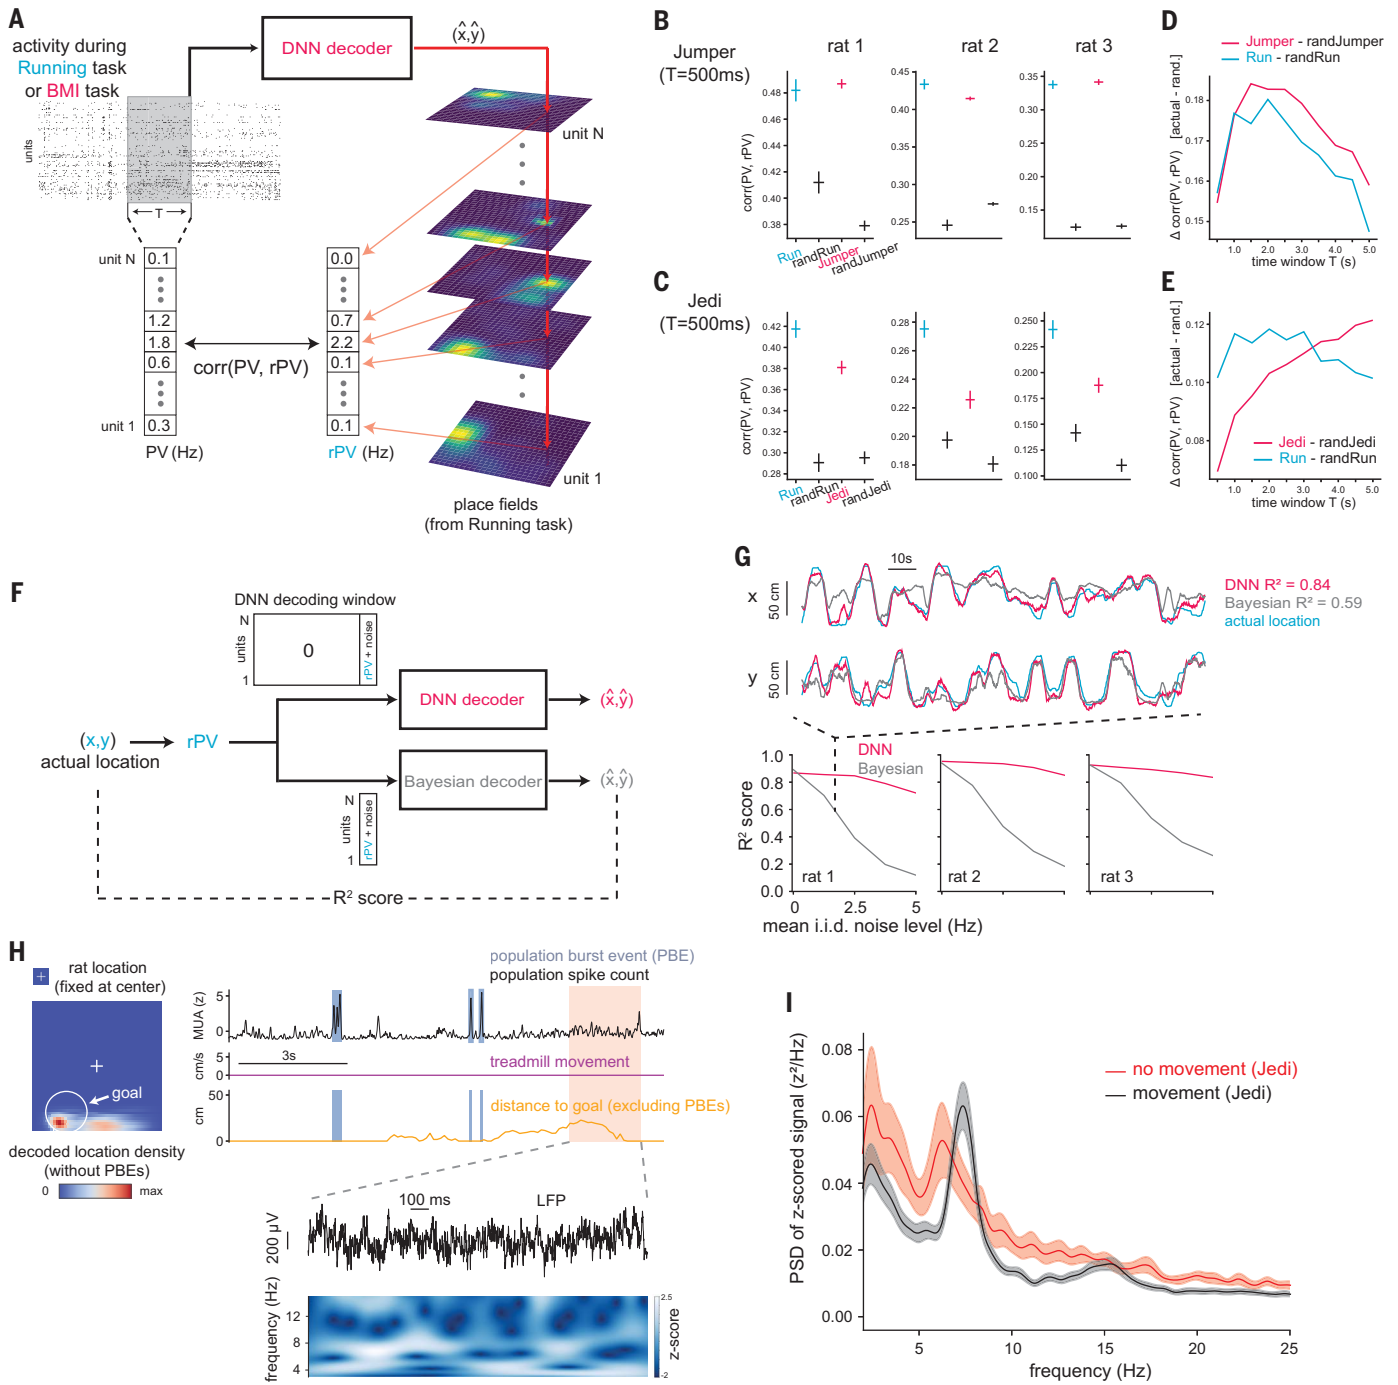


Fig. 4. Volitionally generated nonlocal activity is similar to the activity when the animal is at the corresponding locations and is associated with theta-band power in the LFP. (A to E) The population vector (PV) of ongoing spiking activity was compared with the average place field activity (rPV) at a given location during the running task. (A) Schematic of comparison. (B) Mean correlation of instantaneous (500-ms window) PV during running or jumper task with rPV for the current location (in the running task), current decoded location (in the jumper task), or random location in the running (randRun) or jumper (randJumper) task. (C) Same as (B) but for the Jedi task. For Jedi, only periods when decoded location was near (within 5 cm of) the goal were included [also for (E)]. [(D) and (E)] Correlation of PV with rPV relative to baseline random value as a function of the time integration window for determining the PV.

(F and G) Evaluation of decoder performance when ground truth activity for each location, i.e., the rPV, was input into the decoder. (F) Schematic of evaluation procedure. (G) Comparison of our DNN decoder to Bayesian decoder for different levels of added noise, with example traces using a specific level of noise (top). (H) Distribution of decoded location (left) during Jedi task segment with no treadmill movement (right). Right, from top: summed activity across all units with PBEs identified, treadmill speed, distance of decoded location (excluding data during PBEs) from goal, and close-up of LFP with spectrogram. (I) Power spectral density of z-scored (for pooling across animals) LFP during the Jedi task for periods of treadmill movement and all long segments (≥ 8 s) without treadmill movement. See text and methods in the supplementary materials for details. Here and elsewhere, all confidence intervals (CIs) are 95% CIs.

with the rPVs associated with the decoded locations versus random locations, consistent with the hypothesis. Furthermore, jumper PV-rPV correlations were comparable to running task PV-rPV correlations. In line with this, the example running (Fig. 2B) and jumper (Fig. 2D) trials, which happened to share similar trajectories, showed similar activity patterns across place units over time. PV-rPV correlation scores were, unlike in jumper, lower in Jedi than in the running task using 500-ms windows (Fig. 4C), consistent with noisier generation of nonlocal representations and/or lower firing rates (fig. S13B) in Jedi. However, with longer integration windows (>500 ms) (Fig. 4, D and E), the PVs generated during Jedi matched the rPVs as well as the best match during the running task (note that longer integration times work for Jedi because animals activated goal location representations for extended periods). These results indicate that, during BMI task performance, animals generated nonlocal population activity as similar to the corresponding place field representations as when they actually visited those locations in the running task. Were these place field-like (i.e., rPV-like) patterns what our deep network detected to decode location? While determining what features a deep network uses for decoding is generally not straightforward, inputting a single location's rPV for a brief duration was sufficient to produce accurate location decoding (Fig. 4, F and G), consistent with the decoder being tuned to detect rPV-like activity. In addition, unlike the commonly used Bayesian decoder (45), our decoder was highly robust to noise (Fig. 4G) by design because of the use of data augmentation during training.

Lastly, we analyzed the local field potential (LFP) activity during BMI task performance (Fig. 4, H and I). When animals move, the rodent hippocampal LFP is known to display prominent theta band (~5 to 12 Hz) power, which peaked at ~7.3 Hz during periods of movement in the running and BMI tasks (Fig. 4I). During the extended periods of nonmovement when the animal was performing the Jedi task, the theta peak shifted down to 6.3 Hz (Fig. 4I). Note that, unlike the more continuous theta oscillations during movement, the oscillations during such nonmovement periods tended to be more intermittent.

Discussion

Previous BMI research has yielded major advances in the control of robotic arms, computer cursors, and other devices by activity from the primary motor cortex, premotor cortex, and posterior parietal cortex (24–35). The hippocampal cognitive map has a code that represents space in terms of absolute location in the external environment versus location relative to (e.g., in front of, or to the right or left of) the animal (8–11), and it was unknown

whether a subject could control a BMI by means of this code. With this study, we demonstrated a hippocampal map-based BMI in which the subject is able to control its location or that of other objects by activating location representations in terms of absolute space, independent of where the animal currently is. That is, even though animals generally (but not always) turned their body toward the goal, the activity that needed to be generated differed depending on the location of the goal with respect to the environment. The relatively small amount of training needed for the animals to perform our BMI tasks is in line with our use of a biomimetic decoder (35, 46), that is, one based on the neural code that the subject naturally employs.

In humans, imagining or recalling objects or video clips is accompanied by hippocampal activity in individual neurons similar to that when viewing the original stimuli (47, 48). This suggests that the mechanisms allowing animals to selectively activate their nonlocal hippocampal spatial representations, as we have shown here, could also underlie our ability to actively recall or imagine experiences in other places. The ability of rodents to perform these BMI tasks should thus allow imagination, as well as the voluntary recall of memory, to be investigated using the range of tools available for this model system. More generally, the neural processes engaged here could underlie our capacity to perform “mental time travel”—travel back in time by reexperiencing richly detailed episodic memories and travel forward in time by generating possible future scenarios (49). Mental time travel depends critically on the hippocampus (4–6, 50–52) and enables subjects to internally simulate new experiences according to their world model. This can aid decision-making and facilitate learning in complex situations where trial and error is expensive, as shown using artificial agents (3, 53–55).

Along these lines, the rats in our study could control their hippocampal map-based activity on a timescale of seconds, corresponding to the speed and duration at which humans relive past events or imagine new scenarios. Navigational trajectories each lasted ~10 s, and a virtual object could be held at a remote location for several seconds. This contrasts with the previously described fast (~100 ms) sequences of nonlocal hippocampal activity in awake rodents (i.e., awake replay events, which are associated with population bursts and sharp wave-ripples) thought to be associated with planning (12, 16, 18, 21), and which were not responsible for the performance in our BMI tasks (analysis in which all PBEs were removed). The content of such replay events, which can portray specific routes through the environment starting from the animal's current location, has been shown to be corre-

lated with deliberative (12) and future (16, 18, 21) behavior. However, it is not known whether this content is—or replay content in general can be—under an animal's volitional control. For instance, hippocampal activity displays similar fast sequences during sleep (23), thus nonlocal path generation per se does not appear to require intention. If awake replay is volitionally controlled, these events could represent a brief consideration of alternatives for making a quick decision and be distinct from the more comprehensive mental simulations of possible scenarios that take seconds. Previous work has also described neurons in the hippocampus and related areas whose activity is tuned to the angle to a goal or a salient cue or object relative to the direction the animal is facing (56–59). In addition, hippocampal neurons that are tuned to the location of conspecifics have been found (60, 61). As with fast sequences, whether these forms of activity that reflect locations away from the animal are volitionally controlled is yet to be determined.

Beyond aiding decision-making, the ability to control the content of the hippocampal spatial and episodic memory system could help explain the richness of our inner lives. Finally, the ability to control hippocampal activity to guide oneself or objects to intended locations—and do so with high signal-to-noise readout using our decoder—could lead to new BMI applications for restoring or enhancing function by realizing a subject's high-level intentions with respect to their internal world models.

REFERENCES AND NOTES

- D. L. Schacter, D. R. Addis, R. L. Buckner, *Nat. Rev. Neurosci.* **8**, 657–661 (2007).
- D. L. Schacter et al., *Neuron* **76**, 677–694 (2012).
- J. B. Hamrick, *Curr. Opin. Behav. Sci.* **29**, 8–16 (2019).
- D. Hassabis, D. Kumaran, S. D. Vann, E. A. Maguire, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 1726–1731 (2007).
- D. Hassabis, D. Kumaran, E. A. Maguire, *J. Neurosci.* **27**, 14365–14374 (2007).
- D. R. Addis, A. T. Wong, D. L. Schacter, *Neuropsychologia* **45**, 1363–1377 (2007).
- E. C. Tolman, *Psychol. Rev.* **55**, 189–208 (1948).
- J. O'Keefe, L. Nadel, *The Hippocampus as a Cognitive Map* (Clarendon Press, 1978).
- J. O'Keefe, J. Dostrovsky, *Brain Res.* **34**, 171–175 (1971).
- J. O'Keefe, *Prog. Neurobiol.* **13**, 419–439 (1979).
- M. A. Wilson, B. L. McNaughton, *Science* **261**, 1055–1058 (1993).
- A. Johnson, A. D. Redish, *J. Neurosci.* **27**, 12176–12189 (2007).
- K. Diba, G. Buzsáki, *Nat. Neurosci.* **10**, 1241–1242 (2007).
- T. J. Davidson, F. Kloosterman, M. A. Wilson, *Neuron* **63**, 497–507 (2009).
- M. P. Karlsson, L. M. Frank, *Nat. Neurosci.* **12**, 913–918 (2009).
- B. E. Pfeiffer, D. J. Foster, *Nature* **497**, 74–79 (2013).
- A. M. Wilkenheiser, A. D. Redish, *Nat. Neurosci.* **18**, 289–294 (2015).
- C.-T. Wu, D. Haggerty, C. Kemere, D. Ji, *Nat. Neurosci.* **20**, 571–580 (2017).
- K. Kay et al., *Cell* **180**, 552–567.e25 (2020).
- X. Mou, A. Pokhrel, P. Suresh, D. Ji, *Neuron* **110**, 891–902.e7 (2022).
- J. Widloski, D. J. Foster, *Neuron* **110**, 1547–1558.e8 (2022).
- K. Louie, M. A. Wilson, *Neuron* **29**, 145–156 (2001).
- A. K. Lee, M. A. Wilson, *Neuron* **36**, 1183–1194 (2002).
- J. Wessberg et al., *Nature* **408**, 361–365 (2000).
- J. M. Carmena et al., *PLoS Biol.* **1**, E42 (2003).
- S. Musallam, B. D. Corneil, B. Greger, H. Scherberger, R. A. Andersen, *Science* **305**, 258–262 (2004).
- L. R. Hochberg et al., *Nature* **442**, 164–171 (2006).

28. G. Santhanam, S. I. Ryu, B. M. Yu, A. Afshar, K. V. Shenoy, *Nature* **442**, 195–198 (2006).
29. M. Velliste, S. Perel, M. C. Spalding, A. S. Whitford, A. B. Schwartz, *Nature* **453**, 1098–1101 (2008).
30. P. T. Sadtler *et al.*, *Nature* **512**, 423–426 (2014).
31. T. Aflalo *et al.*, *Science* **348**, 906–910 (2015).
32. G. K. Anumanchipalli, J. Chartier, E. F. Chang, *Nature* **568**, 493–498 (2019).
33. F. R. Willett, D. T. Avansino, L. R. Hochberg, J. M. Henderson, K. V. Shenoy, *Nature* **593**, 249–254 (2021).
34. E. E. Fetz, *J. Physiol.* **579**, 571–579 (2007).
35. M. A. Lebedev, M. A. L. Nicolelis, *Physiol. Rev.* **97**, 767–837 (2017).
36. D. Ishikawa, N. Matsumoto, T. Sakaguchi, N. Matsuki, Y. Ikegaya, *J. Neurosci.* **34**, 5044–5053 (2014).
37. K. Patel, C. N. Katz, S. K. Kalia, M. R. Popovic, T. A. Valiante, *Brain* **144**, 3651–3663 (2021).
38. M. Cerf *et al.*, *Nature* **467**, 1104–1108 (2010).
39. D. Aronov, D. W. Tank, *Neuron* **84**, 442–456 (2014).
40. C. Lai, thesis, University of Cambridge (2020).
41. C. Lai *et al.*, *bioRxiv* 2023.09.14.557854 [Preprint] (2023); <https://doi.org/10.1101/2023.09.14.557854>.
42. C. D. Harvey, F. Collman, D. A. Dombeck, D. W. Tank, *Nature* **461**, 941–946 (2009).
43. G. Chen, J. A. King, Y. Lu, F. Cacucci, N. Burgess, *eLife* **7**, e34789 (2018).
44. J. S. Lee, J. J. Briguglio, J. D. Cohen, S. Romani, A. K. Lee, *Cell* **183**, 620–635.e22 (2020).
45. K. Zhang, I. Ginzburg, B. L. McNaughton, T. J. Sejnowski, *J. Neurophysiol.* **79**, 1017–1044 (1998).
46. K. V. Shenoy, J. M. Carmena, *Neuron* **84**, 665–680 (2014).
47. G. Kreiman, C. Koch, I. Fried, *Nature* **408**, 357–361 (2000).
48. H. Gelbard-Sagiv, R. Mukamel, M. Harel, R. Malach, I. Fried, *Science* **322**, 96–101 (2008).
49. E. Tulving, *Psychologie canadienne* **26**, 1–12 (1985).
50. F. Andelman, D. Hoofien, I. Goldberg, O. Aizenstein, M. Y. Neufeld, *Neurocase* **16**, 426–435 (2010).
51. E. Race, M. M. Keane, M. Verfaellie, *J. Neurosci.* **31**, 10262–10269 (2011).
52. T. Bartsch, J. Döhning, A. Rohr, O. Jansen, G. Deuschl, *Proc. Natl. Acad. Sci. U.S.A.* **108**, 17562–17567 (2011).
53. T. Weber *et al.*, *Adv. Neural Inf. Process. Syst.* **30**, 5690–5701 (2017).
54. D. Ha, J. Schmidhuber, *arXiv:1803.10122 [cs.LG]* (2018).
55. D. Hafner, T. Lillicrap, J. Ba, M. Norouzi, *arXiv:1912.01603 [cs.LG]* (2019).
56. L. Acharya, Z. M. Aghajani, C. Vuong, J. J. Moore, M. R. Mehta, *Cell* **164**, 197–207 (2016).
57. A. Sarel, A. Finkelstein, L. Las, N. Ulanovsky, *Science* **355**, 176–180 (2017).
58. C. Wang *et al.*, *Science* **362**, 945–949 (2018).
59. J. Ormond, J. O’Keefe, *Nature* **607**, 741–746 (2022).
60. T. Danjo, T. Toyozumi, S. Fujisawa, *Science* **359**, 213–218 (2018).
61. D. B. Omer, S. R. Maimon, L. Las, N. Ulanovsky, *Science* **359**, 218–224 (2018).
62. C. Lai, S. Tanaka, T. D. Harris, A. K. Lee, Volitional activation of remote place representations with a hippocampal brain-machine interface, version 1.0, Zenodo (2023); <https://doi.org/10.5281/zenodo.8360872>.

ACKNOWLEDGMENTS

We thank M. Bolstad for work on the virtual reality software; W. Bishop, S. Romani, and X. Zhao for valuable discussions regarding the study; S. Lindo, R. Gattoni, and other members of the Janelia Vivarium team, and B. Lustig, A. Sohn, S. Sawtelle, I. Negrashov, B. Foster, J. Osborne, J. Arnold, R. Rogers, and J. Cohen for technical assistance; and J. Dudman, S. Romani, A. Hantman, T. Wang, and J. Colonell for valuable comments on the manuscript. **Funding:** This work was supported by

the Howard Hughes Medical Institute (A.K.L. and T.D.H.). **Author contributions:** Conceptualization: C.L. and A.K.L. Task design: C.L. BMI and task software and BMI hardware: C.L. Behavioral and recording hardware: S.T. Experiments: S.T. and C.L. Data analysis: C.L. and A.K.L. Writing: C.L., A.K.L., and S.T. Supervision: A.K.L. and T.D.H. **Competing interests:** T.D.H. is also affiliated with the Department of Biomedical Engineering at Johns Hopkins University. **Data and materials availability:** All data needed to assess the findings in this study are publicly available at Zenodo (62). All analysis methods are described in the main text and supplementary materials, and the code for building and applying the decoder is publicly available at Zenodo (62). **License information:** Copyright © 2023 the authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original US government works. <https://www.science.org/about/science-licenses-journal-article-reuse>. This article is subject to HHMI’s Open Access to Publications policy. HHMI lab heads have previously granted a nonexclusive CC BY 4.0 license to the public and a sublicensable license to HHMI in their research articles. Pursuant to those licenses, the author-accepted manuscript (AAM) of this article can be made freely available under a CC BY 4.0 license immediately upon publication.

SUPPLEMENTARY MATERIALS

science.org/doi/10.1126/science.adh5206

Materials and Methods

Figs. S1 to S14

Table S1

References (63–72)

MDAR Reproducibility Checklist

Movie S1

Submitted 9 March 2023; accepted 22 September 2023
10.1126/science.adh5206