PNAS
www.pnas.org

Supplementary Information for

Learning active sensing strategies using a sensory brain-machine interface

Andrew G. Richardson, Yohannes Ghenbot, Xilin Liu, Han Hao, Cole Rinehart, Sam DeLuccia, Solymar Torres Maldonado, Gregory Boyek, Milin Zhang, Firooz Aflatouni, Jan Van der Spiegel, Timothy H. Lucas

Andrew G. Richardson
E-mail: andrew.richardson@uphs.upenn.edu

This PDF file includes:

Other supplementary materials for this manuscript include the following:

Supporting Information Text

Supplementary Results

Alternative encoding site. For the study of active sensing, the sensory BMI paradigm affords flexibility with respect to the encoding location. The principal requirement is that stimulation yield detectable percepts. This requirement can most readily be met by stimulating primary sensory cortical areas (1). To demonstrate an alternative encoding site, stimulating electrodes were implanted in the auditory cortex (A1) of rat Ma (Fig. S2a). After establishing baseline performance with vision, the rat learned over the course of about 100 trials to use ICMS ($\sigma = 45^\circ$, $\mu = 0^\circ$) to locate the platform (Fig. S2b). As with the S1-implanted rats, plateau performance with ICMS (final 30 trials) did not significantly differ from performance with a local visual cue of platform location ($t(71) = .42$, $p = .680$). An example post-learning trial is shown in Movie S4. A pairwise analysis between ICMS trials and catch trials again found that the performance improvement was specifically due to the A1 stimulation (Fig. S2c).

Obstacle task. In many prior studies of ICMS-based rodent navigation, there was negligible delay between the artificial sensory stimulus and the appropriate response and thus little need to memorize the ICMS input (2, 3). In our task, the rats could have used a relatively memory-free strategy of simply swimming toward the stimulus (appropriate for $\mu = 0^\circ$ condition), with immediate turns upon perceived stimulus absence to find the "trail" again. To explore the use of memory, we occasionally placed obstacles in the water maze. On these trials, the rats (Ge and Fr) received ICMS during their initial search of the middle of the tank. If the platform was behind the obstacles (true only for 50% of these trials so that the obstacles themselves were not a visual cue), then the rat was forced to swim around the obstacle, largely in the absence of any further ICMS feedback, to get to the platform. We found that the trained rats were readily able to perform this modified task, circumventing the obstacles when initial feedback indicated the platform was behind them (Movies S5 and S6). This demonstrates that it was possible for the rats to remember ICMS-encoded information to achieve a delayed objective (4). Thus the encoding paradigm achieves a type of hybrid intelligence: providing information about the desired goal while relying on the animal's native intelligence, including their cognitive and memory systems, to overcome any intervening obstacles to reach that goal (5).

4-location task. Our primary results were obtained on a difficult task in which the platform location on each trial was randomized across the entire tank. Previous studies of artificial sensory encoding in rats have used simpler tasks where ICMS is intended to distinguish between just a few discrete choices. To provide a more direct comparison with these studies, we had three rats (Ch, Sn, Ro) perform a simpler task in which the platform could be in one of only four locations (North, South, East, or West), again randomized between trials. The rats were initially trained with local visual cues such that they knew the four possible locations. Then, only ICMS feedback was provided, using ambiguous $\sigma = 90^\circ$ and $\mu = 180^\circ$ encoding parameters. Occasional no-stimulation catch trials were used to provide a performance comparison. An example of eight consecutive trials on this 4-location task is shown in Figure S3a. Performance was qualitatively better with stimulation (trials 2-5, 7-8) than without (trials 1, 6). Performance was quantified with three analyses. First, we analyzed the number of potential platform sites visited (i.e. intersecting the swim path) per trial. For the three rats, the median was 1 or 2 platform-sites/trial for stimulation trials and 3 or 4 platform-sites/trials for no-stimulation trials (Fig. S3b, top row). This result was significant for all three rats (Wilcoxon rank sum tests), as indicated in Figure S3b. Second, we computed the proportion of trials in which the first potential platform site visited was, in fact, the correct one. All three rats were significantly more likely to visit the correct location on the first visit with ICMS feedback than without (chi-squared tests; Fig. S3b, middle row). Third, we computed the proportion of trials in which the first potential platform site visited was the location of the platform on the previous trial. Two of the three rats were significantly more likely to first visit the prior correct location when ICMS was not available compared to when it was available (chi-squared tests; Fig. S3b, bottom row). This final result suggests that the rats often relied on memory, despite its futility given the random sequence of platform locations, when no artificial feedback was provided. Together, the results show that the rats used the low-acuity directional feedback to accomplish the 4-location task, just as they had in the continuous-location task. In addition to providing a point of comparison with prior work, these results show that the rats prioritized the ICMS feedback over the relatively simple alternative of visiting each of the four potential platform sites in turn (i.e. an exhaustive search strategy) (6). This is important as it shows that the encoded sense was not just used out of desperation when an exhaustive search strategy was prohibitive, as it would be in the continuous-location task.

Supplementary Methods

Mathematical model. The observed search strategies, zigzagging and looping, principally involve different yaw maneuvering behaviors. To explore optimality of these strategies, we consider a linear model of yaw inertial dynamics:

$$\begin{aligned} \dot{\theta} &= \omega \\ I\dot{\omega} &= \tau - b\,\omega \end{aligned} \qquad [1]$$

where θ is the heading angle, ω is the angular velocity, τ is the control torque, I is the moment of inertia about the yaw axis, and b is the coefficient of frictional resistance produced by turning in a viscous medium. As described in the Results, the edges of the encoding function, separated by the tuning width parameter σ, contain the most information regarding the goal direction. We interpret zigzagging and looping as strategies to identify these edges. Thus, we consider the objective of encountering both edges with minimum effort (i.e. integrated squared torque) starting from an initial state of $\theta(0) = \omega(0) = 0$ and subject to the

dynamics in Eq. (1). This constrained optimization problem can be solved with the method of Lagrange multipliers using the augmented cost function:

$$J = \boldsymbol{\eta}^{\boldsymbol{T}} \boldsymbol{\Upsilon}(\boldsymbol{x}(t_1)) + \boldsymbol{\nu}^{\boldsymbol{T}} \boldsymbol{\Psi}(\boldsymbol{x}(T)) + \int_0^T \tau^2 dt \tag{2}$$

where $\boldsymbol{x} = [\theta, \omega]^T$ is the state vector, $\boldsymbol{\eta} = [\eta_1, \eta_2]^T$ and $\boldsymbol{\nu} = [\nu_1, \nu_2]^T$ are Lagrange multiplier vectors, and $\boldsymbol{\Upsilon}$ and $\boldsymbol{\Psi}$ provide state constraints at an intermediate time, t_1, when the first edge is encountered and the final time, T, when the second edge is encountered. In particular, the state constraints are:

$$\boldsymbol{\Upsilon} = \begin{bmatrix} \theta(t_1) - \sigma \\ \omega(t_1) - \rho \end{bmatrix}, \ \boldsymbol{\Psi}_{\boldsymbol{zigzag}} = \begin{bmatrix} \theta(T) + \sigma \\ \omega(T) + \rho \end{bmatrix}, \ \boldsymbol{\Psi}_{\boldsymbol{loop}} = \begin{bmatrix} \theta(T) + \sigma - 2\pi \\ \omega(T) - \rho \end{bmatrix} \tag{3}$$

where ρ is the angular speed when each edge is encountered and is left as a free parameter. Notice that the terminal state constraint depends on whether the strategy is to zigzag or loop. Our approach is to minimize Eq. (2) for each case using Pontryagin's minimum principle (7) and then evaluate, for a given parameter set $\boldsymbol{p} = [\sigma, \rho, b, I, T]$, which strategy yields lower cost. For the optimization, we first define the Hamiltonian:

$$H = \tau^2 + \lambda_1 \omega + \lambda_2 \frac{\tau - b\omega}{I} \tag{4}$$

where λ_1 and λ_2 are additional Lagrange multipliers referred to as costate variables. Pontryagin's necessary conditions for minimization of Eq. (2) are:

$$-\dot{\lambda_1} = \frac{\partial H}{\partial \theta} = 0 \tag{5}$$

$$-\dot{\lambda_2} = \frac{\partial H}{\partial \omega} = \lambda_1 - \lambda_2 \frac{b}{I} \tag{6}$$

$$0 = \frac{\partial H}{\partial \tau} = 2\tau + \frac{\lambda_2}{I} \tag{7}$$

Given that λ_1 is constant from Eq. (5), the problem involves solving a system of three ordinary differential equations of the form, $\dot{\hat{\boldsymbol{x}}} = \boldsymbol{A}\hat{\boldsymbol{x}} + \boldsymbol{b}$, governing the dynamics of the state and costate, $\hat{\boldsymbol{x}} = [\theta, \omega, \lambda_2]^T$:

$$\begin{bmatrix} \dot{\theta} \\ \dot{\omega} \\ \dot{\lambda_2} \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & -\frac{b}{I} & -\frac{1}{2I^2} \\ 0 & 0 & \frac{b}{I} \end{bmatrix} \begin{bmatrix} \theta \\ \omega \\ \lambda_2 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ -\lambda_1 \end{bmatrix} \tag{8}$$

The solution to this linear system is:

$$\hat{\boldsymbol{x}}(t) = \boldsymbol{\Phi}\boldsymbol{c} + \boldsymbol{\Phi} \int \left(\boldsymbol{\Phi}^{-1} \boldsymbol{b} \right) dt \tag{9}$$

where $\boldsymbol{\Phi} = e^{\boldsymbol{A}t}$ is the state transition matrix and the constant vector, $\boldsymbol{c}$, is determined by the boundary conditions. However, $\boldsymbol{c}$, is determined separately for $0 < t < t_1$ and $t_1 < t < T$ to turn the three-point boundary value problem into a pair of two-point boundary value problems (8). The additional required continuity constraints just before ($^-$) and after ($^+$) $t = t_1$ are:

$$\lambda_1^-(t_1) = \lambda_1^+(t_1) + \eta_1 \tag{10}$$

$$\lambda_2^-(t_1) = \lambda_2^+(t_1) + \eta_2 \tag{11}$$

$$H^-(t_1) = H^+(t_1) \tag{12}$$

$$\theta^-(t_1) = \theta^+(t_1) \tag{13}$$

$$\omega^-(t_1) = \omega^+(t_1) \tag{14}$$

Thus, the costate variables are discontinuous at t_1 while the state variables are continuous. The optimal time of interception of the first edge, t_1^*, is determined by Eq. (12). This value is then substituted into Eq. (9) to find the optimal state-costate trajectory, $\hat{\boldsymbol{x}}^*$. Finally, the optimal torque, $\tau^*(t)$, is determined from Eq. (7). For each parameter set, $\boldsymbol{p}$, we determine the optimal strategy by comparing the minimum effort corresponding to the two different terminal state constraints in Eq. (3):

$$\int_0^T \tau_{zigzag}^{*2} > \int_0^T \tau_{loop}^{*2} \Rightarrow \text{looping is optimal} \tag{15}$$

Model parameters. The optimal strategy was determined for a range of parameter values estimated from the empirical data. First, I and b were calculated from the rats' morphology. For the former, we approximated the rat body as an ellipsoid. For the latter, we used the formula given in (9):

$$I = \frac{1}{5} m \left[\left(\frac{L}{2} \right)^2 + \left(\frac{D}{2} \right)^2 \right] \quad [16]$$

$$b = \frac{\pi \mu L^3}{3 \, ln \left(\frac{L}{2D} \right)} \quad [17]$$

where $m = 0.3$ kg is the body weight, $L = 20$ cm is the body length, $D = 6$ cm is the body diameter, and $\mu = 8.9$ x 10^{-4} kg/(m·s) is the dynamic viscosity of water at 25°C. The ratio, $b/I = 0.023$ s^{-1}, is the inverse of the time constant of the dynamics in Eq. (1) (9). Second, we calculated the distribution of angular swim speeds for each rat. The average 95th percentile of these distributions, 269 ± 71 deg/s, was used to define the upper limit of ρ. Third, observed values of T (i.e. total time needed to reach both edges of the encoding function) varied with σ but were typically in the range of 0.3 to 0.8 s.

Electrode localization. At the conclusion of the experiment, each rat was anesthetized as in the implantation procedure and transcardially perfused with heparinized saline followed by 10% buffered formalin. After 24 h in formalin, the brain was extracted from the skull and a coronal block of tissue around the electrode implant site was washed, dehydrated, and embedded in paraffin wax. Serial sections (8 μm thickness) of this block were stained with hematoxylin and eosin. The electrodes were localized through comparison of the sections to a brain atlas (10). Results are shown in Figure S4. In all cases, the electrodes were judged to be within the cortical target (S1 or A1). However, there was heterogeneity in the exact S1 representation stimulated. The electrodes in rats Fr, Ge, Sn, and Ch were localized to the S1 barrel field, while in rats Sa and Ro they were in the more anterior-medial S1 forelimb and hindlimb representations, respectively. This anatomical data agrees with the functional data obtained upon stimulating these areas to determine perceptual thresholds. In all four rats with electrodes in the barrel fields, we observed clear twitching or scratching of the contralateral whisker pad in response to supraliminal ICMS. In the two rats with electrodes in limb representations, evoked responses were less obvious (subtle contralateral turning) and were not directed at the whiskers.

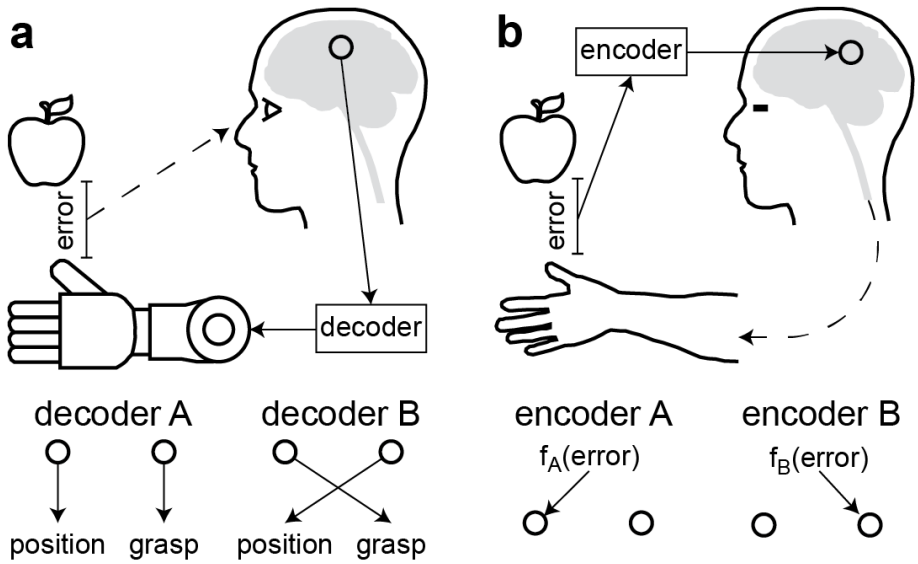

Fig. S1. BMI paradigms to study learning. (a) Illustration of a motor BMI for learning procedural tasks. Learning occurs after changing the relationship of how recorded neural activity controls the effector (i.e. switching from decoder A to decoder B). **(b)** Illustration of a sensory BMI for learning active sensing tasks. Learning occurs after changing the relationship between the task relevant variable (e.g. error) and the neural stimulation. In both illustrations, dashed and solid lines indicate natural and artificial pathways, respectively.

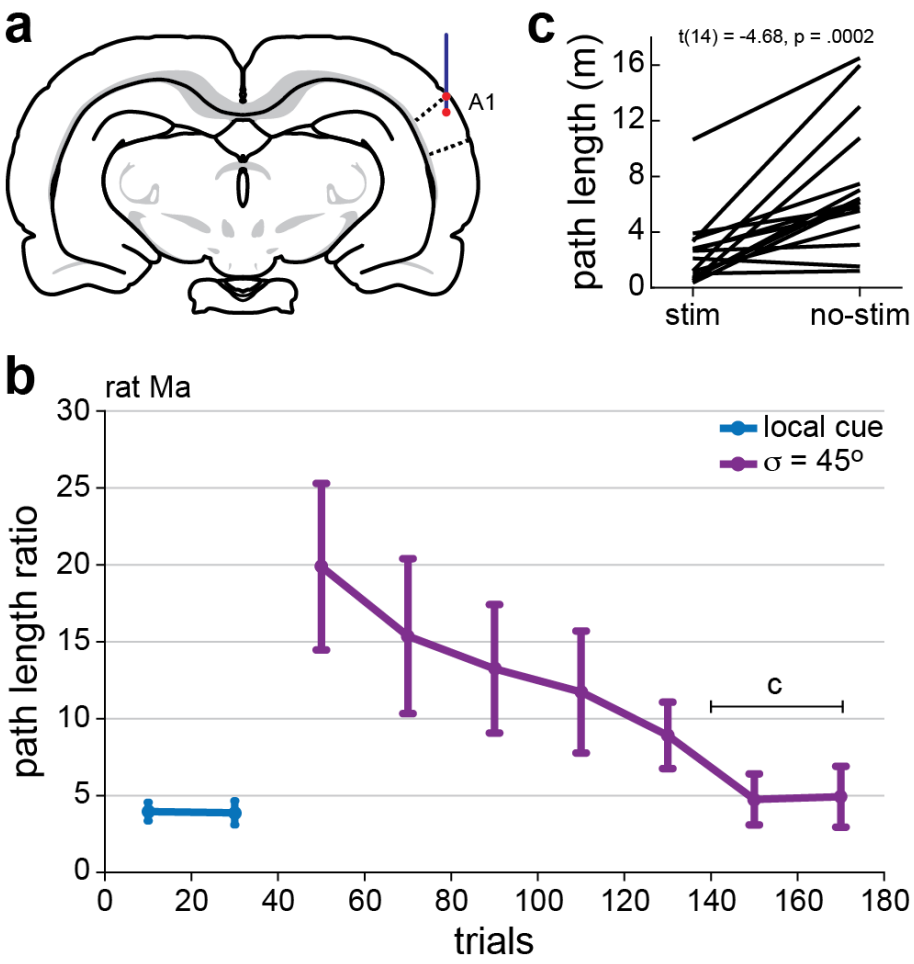

Fig. S2. Learning with high acuity ICMS feedback delivered to A1. (a) Illustration (11) of the bipolar stimulation sites (red circles) in rat Ma targeting primary auditory cortex (A1). **(b)** Task performance (mean ± 95% c.i. for sequential groups of 20 trials) with a local visual cue (dark blue) and with ICMS (purple). Encoding parameters for the latter were $\sigma = 45^{\circ}$ and $\mu = 0^{\circ}$. **(c)** Results of a paired statistical analysis of swim path length with and without ICMS on trials with matched platform locations. The statistical test was one-tailed. No-stimulation catch trials occurred during the post-learning trial blocks indicated in **(b)**.

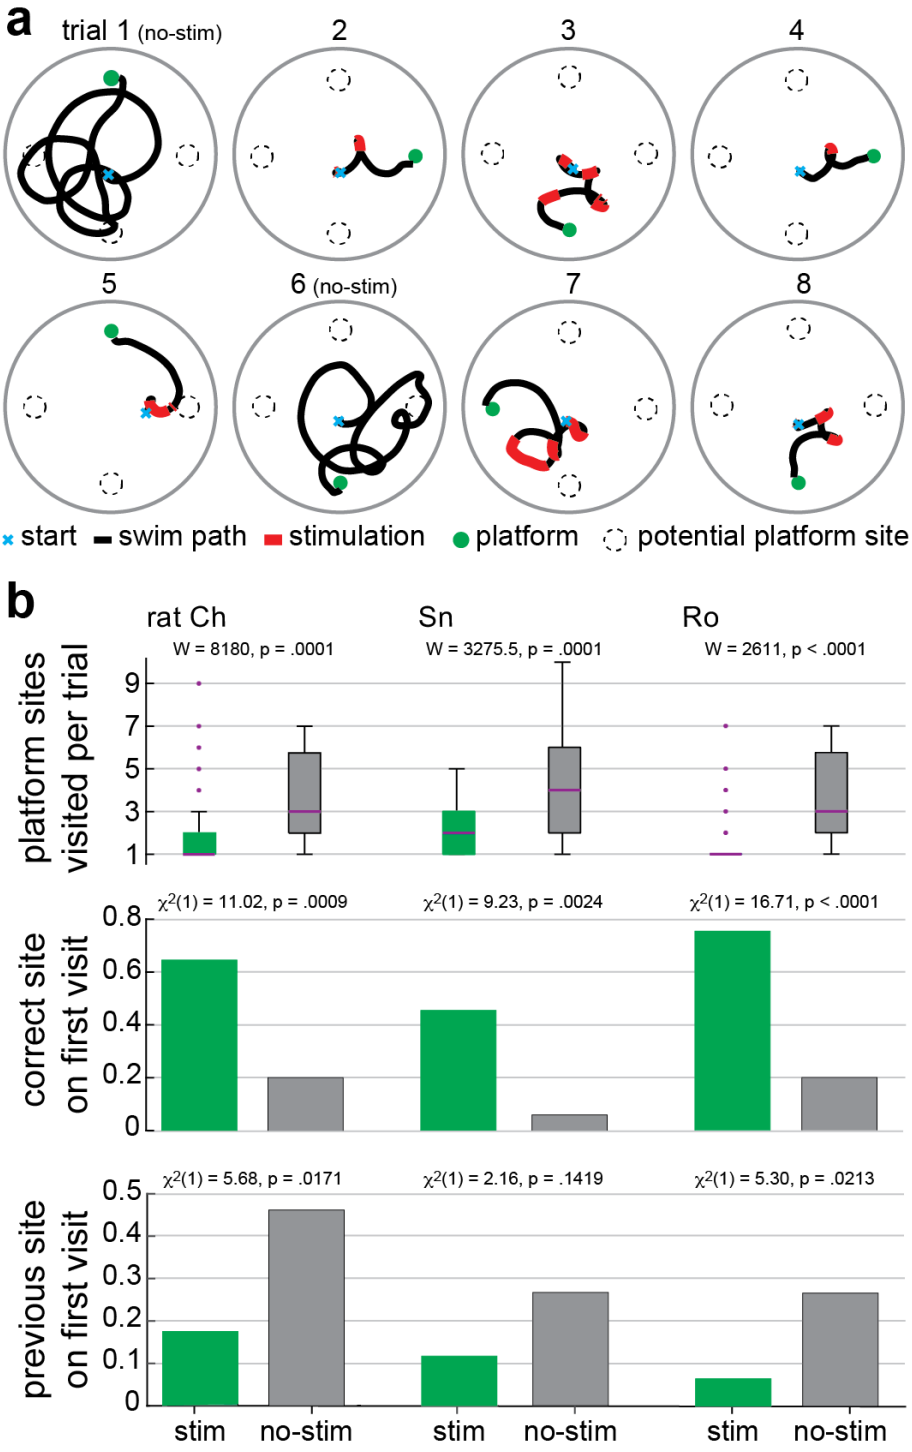

Fig. S3. Prioritizing low acuity ICMS feedback on an easier task. (a) An example of eight consecutive trials from rat Ro demonstrating the use of goal-direction feedback on an easier task with only four potential platform locations. In this example, trials 1 and 6 were catch trials in which no stimulation was delivered. **(b)** Statistical analysis comparing performance on the 4-location task in the stimulation ($\sigma = 90°$, $\mu = 180°$) and no-stimulation (catch trial) conditions for three rats. **Top row,** Comparison of the number of potential platform sites visited (i.e. intersecting the swim path) per trial. Wilcoxon rank sum test reported. **Middle row,** Comparison of the proportion of trials in which the first potential platform site visited was, in fact, the correct site (i.e. the actual platform location). **Bottom row,** Comparison of the proportion of trials in which the first potential platform site visited was the location of the platform on the previous trial.

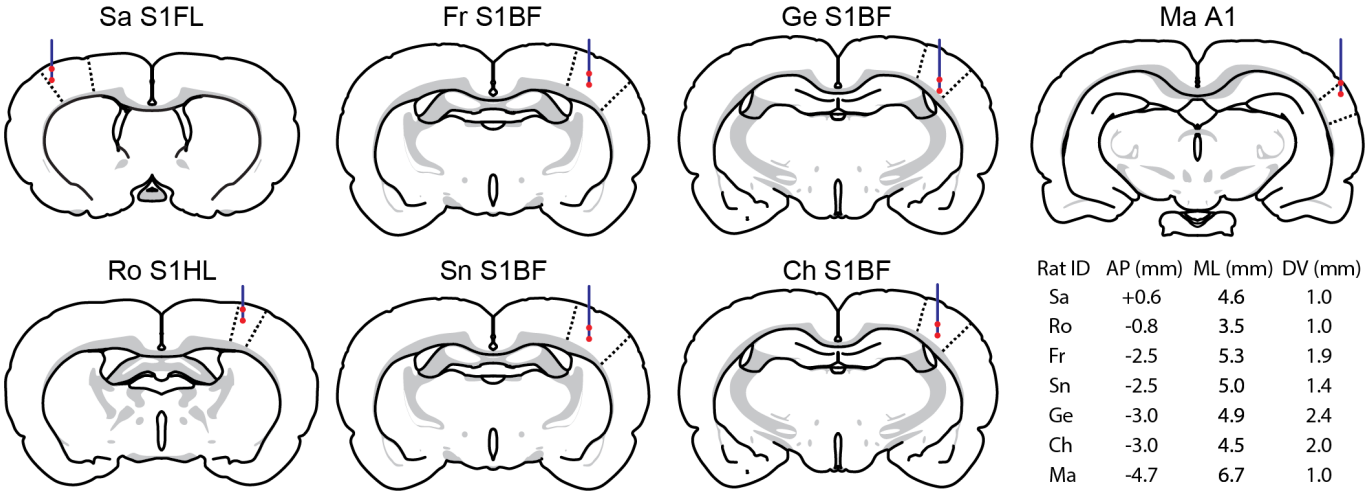

Rat ID	AP (mm)	ML (mm)	DV (mm)
Sa	+0.6	4.6	1.0
Ro	-0.8	3.5	1.0
Fr	-2.5	5.3	1.9
Sn	-2.5	5.0	1.4
Ge	-3.0	4.9	2.4
Ch	-3.0	4.5	2.0
Ma	-4.7	6.7	1.0

Fig. S4. Electrode localization. Illustrations (11) of the bipolar electrode location (blue line with red circles representing the electodes) in each rat, as determined by histology. S1FL = forelimb representation of primary somatosensory cortex, S1HL = hindlimb representation of primary somatosensory cortex, S1BF = barrel field (vibrissal representation) of primary somatosensory cortex, A1 = primary auditory cortex. The table provides the estimated anterior-posterior (AP), medial-lateral (ML), and dorsal-ventral (DV) coordinates of the distal electrode. AP, ML, and DV coordinates are relative to bregma, midline, and the pia, respectively.

Table S1. Task and encoding conditions for each subject

				Encoding functions	
Rat	Trials §	Days †	Tasks ‡	σ	μ
Ge	609	52	CL, OB	5, 15, 45, 90	0
Fr	337	33	CL, OB	5, 15, 45	0
Sa	394	31	CL	90	0, 180
Ma	202	20	CL	45	0
Sn	469	45	4L, CL	90	180
Ro	420	33	4L, CL	90	180
Ch	259	21	4L	90	180

§ total number of trials, including visual cue, ICMS, and catch trials

† number of daily sessions over which trials were conducted

‡ continuous-location (CL) task, 4-location (4L) task, obstacle (OB) task

Movies

Movie S1. Pre-learning trial. Rat Fr, $\sigma = 45°$, $\mu = 0°$.

Movie S2. Post-learning trial. Rat Fr, $\sigma = 45°$, $\mu = 0°$.

Movie S3. Upper acuity limit. Rat Ge, $\sigma = 5°$, $\mu = 0°$.

Movie S4. ICMS of A1. Rat Ma, $\sigma = 45°$, $\mu = 0°$.

Movie S5. Obstacle trial. Rat Ge, $\sigma = 15°$, $\mu = 0°$.

Movie S6. Obstacle trial. Rat Fr, $\sigma = 15°$, $\mu = 0°$.

Movie S7. Low acuity feedback. Rat Ch, $\sigma = 90°$, $\mu = 180°$.

Movie S8. Low acuity feedback. Rat Ro, $\sigma = 90°$, $\mu = 180°$.

References

1. Histed MH, Ni AM, & Maunsell JH (2013) Insights into cortical mechanisms of behavior from microstimulation experiments. Prog Neurobiol 103:115-130.
2. Talwar SK, et al. (2002) Rat navigation guided by remote control. Nature 417(6884):37-38.
3. Lee MG, et al. (2010) Operant conditioning of rat navigation using electrical stimulation for directional cues and rewards. Behav Processes 84(3):715-720.
4. Romo R, Hernandez A, Zainos A, Brody CD, & Lemus L (2000) Sensing without touching: psychophysical performance based on cortical microstimulation. Neuron 26(1):273-278.
5. Gao L, et al. (2013) Ratbot automatic navigation by electrical reward stimulation based on distance measurement in unknown environments. Conf Proc IEEE Eng Med Biol Soc 2013:5315-5318.
6. Baldi E, Lorenzini CA, & Corrado B (2003) Task solving by procedural strategies in the Morris water maze. Physiol Behav 78(4-5):785-793.
7. Pontryagin LS, Boltyanskii VG, Gamkrelidze RV, & Mishechenko EF (1962) The mathematical theory of optimal processes. Interscience Publishers Inc., New York.
8. Bryson AE & Ho YC (1975) Applied optimal control. Taylor & Francis Group, New York.
9. Fry SN, Sayaman R, & Dickinson MH (2003) The aerodynamics of free-flight maneuvers in drosophila. Science 300:495-498.
10. Paxinos G, Watson C (2004) The rat brain in stereotaxic coordinates (5th edition). Academic Press, New York.
11. Swanson LW (2004) Brain maps: structure of the rat brain (3rd edition). Elsevier, Amsterdam.