Article

Motor learning refines thalamic influence on motor cortex

https://doi.org/10.1038/s41586-025-08962-8

Received: 10 May 2024

Accepted: 1 April 2025

Published online: 07 May 2025

Check for updates

Assaf Ramot[1,2,3,4,7], Felix H. Taschbach[1,5,7], Yun C. Yang[1,2,3,4], Yuxin Hu[1,2,3,4], Qiyu Chen[1,2,3,4], Bobbie C. Morales[1,2,3,4], Xinyi C. Wang[1,2,3,4], An Wu[1,2,3,4,8], Kay M. Tye[5,6], Marcus K. Benna[1]✉ & Takaki Komiyama[1,2,3,4,6]✉

The primary motor cortex (M1) is central for the learning and execution of dexterous motor skills[1–3], and its superficial layer (layers 2 and 3; hereafter, L2/3) is a key locus of learning-related plasticity[1,4–6]. It remains unknown how motor learning shapes the way in which upstream regions activate M1 circuits to execute learned movements. Here, using longitudinal axonal imaging of the main inputs to M1 L2/3 in mice, we show that the motor thalamus is the key input source that encodes learned movements in experts (animals trained for two weeks). We then use optogenetics to identify the subset of M1 L2/3 neurons that are strongly driven by thalamic inputs before and after learning. We find that the thalamic influence on M1 changes with learning, such that the motor thalamus preferentially activates the M1 neurons that encode learned movements in experts. Inactivation of the thalamic inputs to M1 in experts impairs learned movements. Our study shows that motor learning reshapes the thalamic influence on M1 to enable the reliable execution of learned movements.

Reliable execution of learned movements is a fundamental aspect of adaptive behaviour, and is essential for an animal's survival and well-being. During repeated practice of a motor skill, the variability of movements across trials decreases and the speed at which a desired outcome is achieved increases, in a process known as motor learning. During this process, the motor circuits in the brain undergo changes that facilitate the reliable and efficient execution of the learned motor skill[7,8].

The primary motor cortex (M1) is a central locus for motor learning and execution in the mammalian brain[1–3]. Even though not all movements require M1 (ref. 9), and some M1-dependent motor skills can become independent of M1 after long-term training[10,11], M1 is clearly essential for the initial learning of many dexterous motor skills. In particular, the superficial layer L2/3 of M1 is a major locus of changes during motor learning. Synapses onto M1 L2/3 excitatory neurons reorganize during motor learning[1,4–6], and this coincides with the emergence of reproducible spatio-temporal activity in the M1 L2/3 neural population that accompanies learned movements[1]. M1 L2/3 neurons drive deeper layers of M1 that house neurons projecting subcortically and serve as the output layer of M1 (ref. 12). A model has emerged in which M1 L2/3 receives long-range inputs that interact with local recurrent circuits to generate the skill-specific ensemble activity that dynamically drives deeper-layer neurons to execute the learned motor skill. However, several key questions remain, including the identity of the brain area that provides the key input to drive the skill-specific M1 L2/3 activity, the specific M1 L2/3 neurons it activates, and the way this input pathway is shaped during motor learning. Cellular-level interactions across brain areas are poorly understood, owing partially to the technical difficulties of identifying effective connectivity across areas—especially longitudinally throughout learning.

Here, we investigate the main sources of long-range inputs to M1 L2/3. Using longitudinal axonal calcium imaging, we identify the motor thalamus as the input area that provides the strongest excitation to M1 before and during a learned movement. We then establish a methodology to identify the sparse group of M1 L2/3 neurons that are strongly driven by thalamic inputs, and characterize their functional properties during behaviour. We develop a computational method for aligning paired high-dimensional multimodal data across individuals, and use it to uncover the unique encoding properties of the thalamus-excited M1 L2/3 neurons. A longitudinal analysis reveals that learning refines the thalamic influence on M1, such that thalamic inputs strongly activate movement-preceding neurons in experts. We propose that the precise reorganization of the thalamocortical pathway is a crucial component of motor learning.

Motor learning and M1 activity

To investigate M1 circuits that are associated with learned movements, we used a cued lever-press task[1]. In this task, water-restricted mice were trained under head fixation daily for one session per day for two weeks. During training sessions, mice used their left forepaw to grasp a lever. An auditory cue signalled the answer period during which a lever press past the threshold produced a water reward (Fig. 1a). In a comparison of the expert stage (days 13–14) and the beginner stage (days 1–2), we found that the number of rewarded trials increased (Fig. 1b) in the expert stage, and the times from cue to movement onset and from cue to reward decreased (Fig. 1c). Movement trajectories became more stereotyped in the expert stage (Fig. 1d and Extended Data Fig. 1), a hallmark of motor learning.

[1]Department of Neurobiology, University of California San Diego, La Jolla, CA, USA. [2]Center for Neural Circuits and Behavior, University of California San Diego, La Jolla, CA, USA. [3]Department of Neurosciences, University of California San Diego, La Jolla, CA, USA. [4]Halıcıoğlu Data Science Institute, University of California San Diego, La Jolla, CA, USA. [5]Salk Institute for Biological Studies, Howard Hughes Medical Institute, La Jolla, CA, USA. [6]Kavli Institute for the Brain and Mind, La Jolla, CA, USA. [7]These authors contributed equally: Assaf Ramot, Felix H. Taschbach. [8]Deceased: An Wu. ✉e-mail: mbenna@ucsd.edu; tkomiyama@ucsd.edu

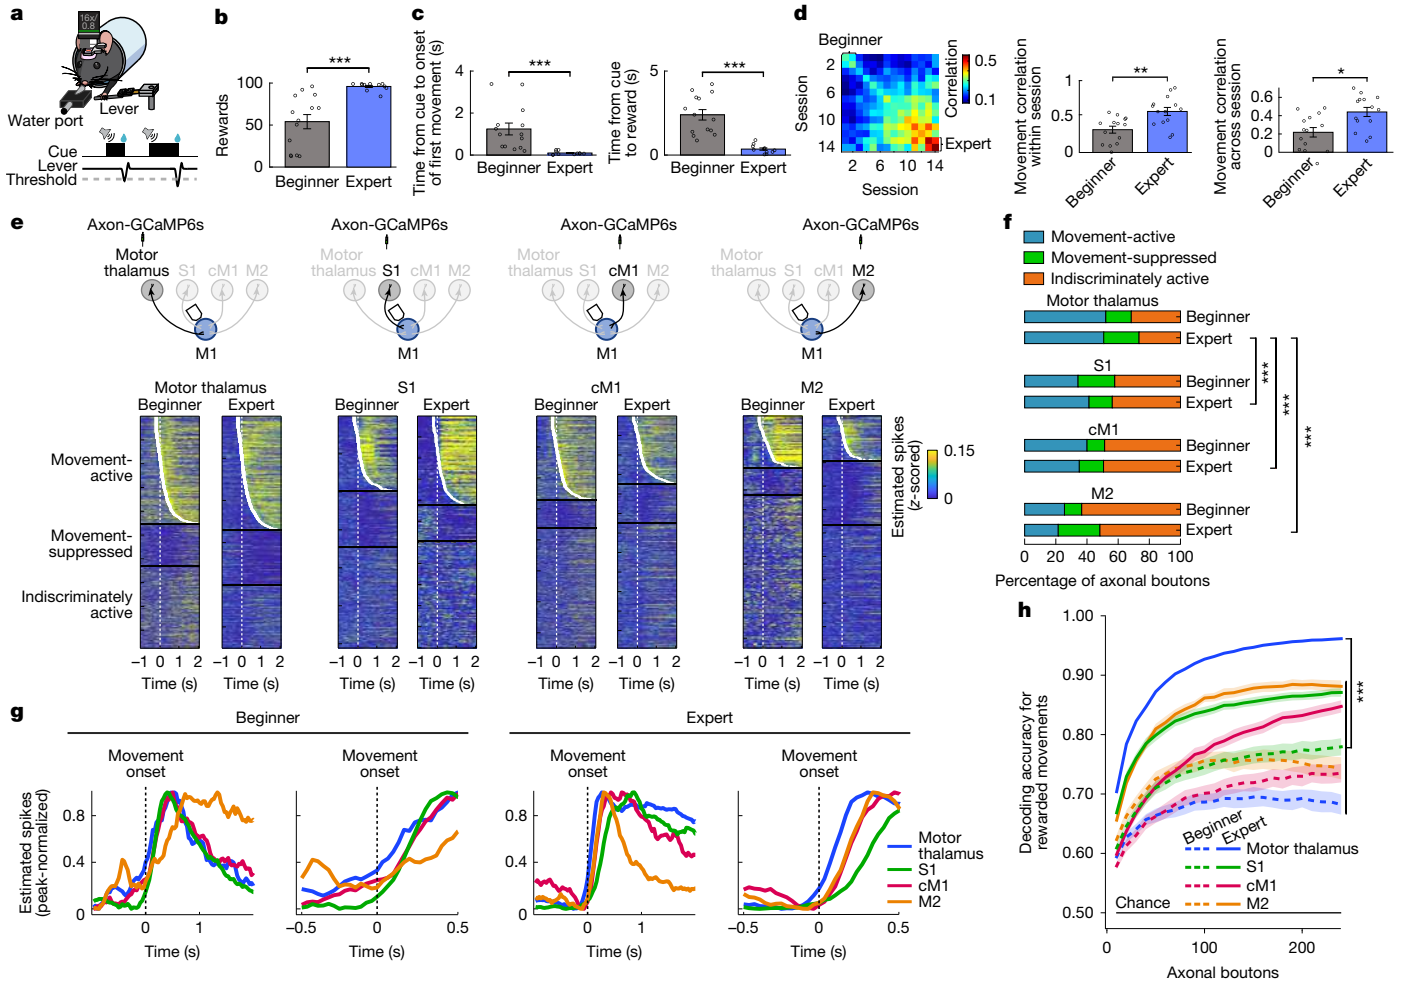

Fig. 1 | Thalamic inputs to M1 preferentially encode learned movements in experts. a, Schematic of the experimental set-up and task structure. Adapted from ref. 4. **b**, The number of rewarded trials increased over training ($P < 0.001$, Wilcoxon rank-sum test). Circles correspond to individual sessions. **c**, The time from cue onset to movement onset and from cue onset to reward decreased over training ($P < 0.001$, rank-sum test). **d**, Left, median trial-by-trial correlation coefficients of rewarded movement trajectories, averaged across mice. Rewarded movements became more similar within (middle) and across (right) sessions ($P < 0.01$ and $P < 0.05$, rank-sum test, respectively). **e**, Top, schematic of injections to express axon-GCaMP6s in different input areas and image their axons innervating M1. Bottom, trial-averaged activity of movement-active (top, sorted by onset timing, white), movement-suppressed (middle) and indiscriminately active (bottom) axonal boutons in thalamic ($n = 7$ mice), S1 ($n = 4$ mice), cM1 ($n = 7$ mice) and M2 ($n = 7$ mice) inputs to M1 during beginner (left) and expert (right) sessions, aligned to the onset of rewarded movements (dashed lines). Each row represents one axonal bouton. **f**, The fractions of motor thalamus, S1, cM1 and M2 axonal boutons that are movement-active, movement-suppressed and indiscriminately active. *n* for movement-active, movement-suppressed and indiscriminately active axonal boutons (beginner/expert), respectively: thalamus: 1,133/1,441, 356/644 and 687/755; S1: 432/238, 297/86 and 532/252; cM1: 638/454, 180/201 and 776/639; and M2: 166/219, 72/272 and 412/525. In experts, motor thalamic inputs exhibited the highest fraction of movement-active boutons ($\chi^2 = 288.736$, degrees of freedom = 3, $P < 0.001$). Post-hoc pairwise *z*-tests with false-discovery-rate correction confirmed significant differences between motor thalamic inputs and all other groups ($P < 0.001$ for each comparison). **g**, Population-average activity of movement-active inputs during the beginner (left) and expert (right) sessions, aligned to the onset of rewarded movements (peak-normalized). In experts, the population-average activity onset of only thalamic inputs preceded movements. **h**, Cross-validated decoding accuracy for distinguishing time points for rewarded movements versus no movements. The activity across multiple mice was combined as pseudo-simultaneous populations (Methods). The decoding accuracy for thalamic inputs in experts was significantly higher than that for all other inputs (pairwise *t*-tests with Holm–Bonferroni multiple-comparison correction, all corrected $P < 10^{-7}$). Shaded areas show 95% confidence intervals of the mean over cross-validation folds. For panels **b**–**d**, **f** and **h**, $^*P < 0.05$, $^{**}P < 0.01$ and $^{***}P < 0.001$.

We next examined the ensemble activity of L2/3 neurons in the caudal forelimb area of the right M1 by performing longitudinal two-photon calcium imaging of the same population of neurons across training ($n = 10$ imaging fields from 7 mice) (Extended Data Fig. 2a). A substantial proportion of neurons exhibited movement-related activity, categorized as either 'movement-active' or 'movement-suppressed' (Extended Data Fig. 2b–d), with their activity tiling the duration of rewarded movements (Extended Data Fig. 1). The trial-by-trial correlation of population activity patterns during rewarded movements increased during learning, indicating the emergence of a reproducible activity pattern (Extended Data Fig. 2e).

Thalamic inputs encode learned movements

Previous studies showed that M1 activity is essential for executing learned movements after two weeks of training[1,10,11]. M1 receives inputs from various brain regions[13–15]. We hypothesized that specific inputs initiate the dynamics of M1 L2/3 to drive the learned movement. To test this idea, we started by identifying the primary sources of long-range inputs to M1 L2/3 neurons. We used a published dataset from an experiment in which rabies-virus-based monosynaptic retrograde labelling was performed with M1 L2/3 excitatory neurons as starter cells[16]. We identified the ipsilateral motor thalamus (ventral anterior-lateral complex of

the thalamus (VAL) and ventral medial nucleus of the thalamus (VM)) and surrounding thalamic nuclei (ventral posterolateral nucleus of the thalamus (VPL), ventral posteromedial nucleus of the thalamus (VPM) and posterior complex of the thalamus (PO)), the ipsilateral somatosensory cortex (S1), the contralateral M1 (cM1) and the ipsilateral secondary motor cortex (M2) as the four main input areas that accounted for around 88% of the long-range inputs to M1 L2/3 (Extended Data Fig. 3). Next we examined the activity of these long-range inputs to the right M1 during motor learning. To do this, we expressed the axon-targeted calcium sensor axon-GCaMP6s[17] in one of the input areas in each mouse, and imaged the activity of their axons in M1 during training (Fig. 1e). By analysing fluorescence from individual axonal boutons, we found that all four inputs carry substantial movement-related activity throughout learning (Fig. 1f and Extended Data Fig. 4), but thalamic inputs exhibited the highest proportion of movement-active axonal boutons (Fig. 1f). Especially in expert sessions, the movement-active thalamic inputs started before the onset of the learned movement, whereas the other inputs were relatively delayed, indicating that thalamic inputs could contribute to the initiation of the learned movement (Fig. 1g, right). Next, to quantify the amount of movement-related activity in each input, we used their activity to discriminate the time points of rewarded movements versus time points without movement, using linear SVM classifiers. Consistent with the other analysis above, thalamic inputs at the expert sessions outperformed the other inputs' decoding of rewarded movements (Fig. 1h). Thus, after training, learned movements are preferentially encoded by thalamic inputs to M1. Together, these results suggest that thalamic inputs have a key role in exciting M1 L2/3 to drive learned movements.

Identifying thalamus-excited M1 neurons

On the basis of these observations, we hypothesized that thalamocortical inputs activate a specific subset of M1 L2/3 neurons to drive the dynamics of the M1 L2/3 population underlying the execution of the learned movement. To test this idea, we sought to identify the M1 L2/3 neurons strongly driven by thalamic inputs, using optogenetic stimulation of thalamic axons in M1. To this end, we first validated our ability to reliably excite thalamic axons in M1. We injected Cre-dependent AAV encoding ChrimsonR[18] and GCaMP6f[19] into the motor thalamus of Vglut2-Cre mice[20], and imaged thalamocortical axons in M1 (Extended Data Fig. 5a). These axons showed strong and reliable activity in response to optogenetic stimulation (Extended Data Fig. 5b). After this validation, we combined optogenetic stimulation of thalamic axons with two-photon calcium imaging of M1 L2/3 neurons by expressing ChrimsonR in the motor thalamus of CaMKII-tTA::tetO-GCaMP6s::Vglut2-Cre triple transgenic mice, in which GCaMP6s is expressed in cortical excitatory neurons (Fig. 2a–c). These optogenetic mapping (opto-mapping) sessions were done before (session 0) and after (session 15) the two-week training (Fig. 2a). During the opto-mapping sessions, we stimulated thalamic axons in M1 while we imaged the activity of M1 L2/3 neurons. This approach allowed us to identify the M1 L2/3 neurons that respond to the stimulation of thalamic inputs (Fig. 2d and Extended Data Fig. 6). A subset (around 9%) of M1 L2/3 neurons showed significantly increased activity during thalamic stimulation ('Th-excited'; we note that they might not all be monosynaptically driven by thalamic inputs). A fraction of neurons showed increased activity after the offset of thalamic stimulation ('Th-rebound'), but these neurons were very rare (around 1%) and we did not study them further (Fig. 2e). We refer to all other neurons in the field of view (FOV) as 'Th-non-responsive'.

We next considered the possibility that stimulation of thalamic inputs leads to the activation of Th-excited neurons through indirect stimulation of other brain areas (that is, stimulation of axons in M1 activates thalamic neurons through backpropagation of action potentials, which in turn activate other brain areas, which then indirectly excite M1 neurons). To address this issue, we inactivated the motor thalamus by local administration of the GABAA agonist muscimol and stimulated thalamic axons in M1 while imaging the activity of M1 L2/3 neurons (Extended Data Fig. 7a). The number of excited neurons was similar before and after the injection of muscimol into the motor thalamus (Extended Data Fig. 7b,c). Even though the thalamic injection of muscimol might not eliminate all collateral activations, these results support the notion that Th-excited neurons in M1 are excited owing to the local effects of thalamic axons in M1.

Although the overall fraction of Th-excited neurons was similar in sessions 0 and 15 (Fig. 2e), the identity of Th-excited neurons changed substantially. A subset remained consistently Th-excited before and after learning ('stable' Th-excited; around half), but others either gained ('gain' Th-excited neurons) or lost ('loss' Th-excited neurons) this property over the course of training (Fig. 2f).

Th-excited neurons encode learned movements

Equipped with the identity of M1 L2/3 neurons that are strongly driven by thalamocortical inputs, we asked whether their activity patterns during training are distinct from those of Th-non-responsive neurons. We defined Th-excited neurons in beginner and expert stages separately (session 0 for beginner mice and session 15 for expert mice; Fig. 2a). We examined whether Th-excited and Th-non-responsive populations differ in their overall movement-related information using the same decoding analysis as in Fig. 1. In beginners, Th-excited neurons slightly outperformed Th-non-responsive neurons when the number of neurons was matched. Furthermore, Th-excited neurons showed a strong increase in their movement encoding during learning, and Th-excited neurons in expert mice contained substantially more movement-related information than did the Th-non-responsive population (Fig. 2g). This strong encoding of rewarded movements in expert Th-excited neurons aligns with our similar observation in the activity of thalamic inputs (Fig. 1h).

We proceeded to evaluate the movement-related activity in Th-excited and Th-non-responsive populations, and found that Th-excited neurons exhibited higher activity than Th-non-responsive neurons near the initiation of rewarded movements, especially in experts (Fig. 2h,i and Extended Data Fig. 8). This higher movement-related activity of Th-excited neurons is due partially to the higher fraction of movement-excited neurons in the Th-excited population (Fig. 2j). The fraction of movement-excited neurons in the Th-excited population increased during learning, and in expert mice, a majority of Th-excited neurons showed movement-active responses (Fig. 2j, bottom).

The strong encoding of rewarded movements by Th-excited neurons was not only due to the larger fraction of movement-active neurons. A decoding analysis using only movement-active neurons found that Th-excited movement-active neurons in experts are more strongly modulated by rewarded movements than are Th-non-responsive movement-active neurons. This contrasts with similar decoding accuracies between movement-active neurons in Th-excited versus Th-non-responsive populations in beginners (Fig. 2k). Consistent with these observations, movement-aligned population-average activity showed stronger movement modulation of Th-excited movement-active neurons in experts (Fig. 2m). This difference was not observed in beginners (Fig. 2l). Th-excited movement-active neurons were active in a larger fraction of rewarded movements (higher 'reliability') than Th-non-responsive movement-active neurons, during expert sessions only (Fig. 2n, bottom). These results suggest that in experts, Th-excited neurons exhibit stronger movement-related modulation than Th-non-responsive neurons. However, we acknowledge that strictly separating the effects of prevalence and individual response strength is difficult in the presence of noise.

In addition, the timing of activity onset relative to the onset of rewarded movements showed learning-related changes. In beginners,

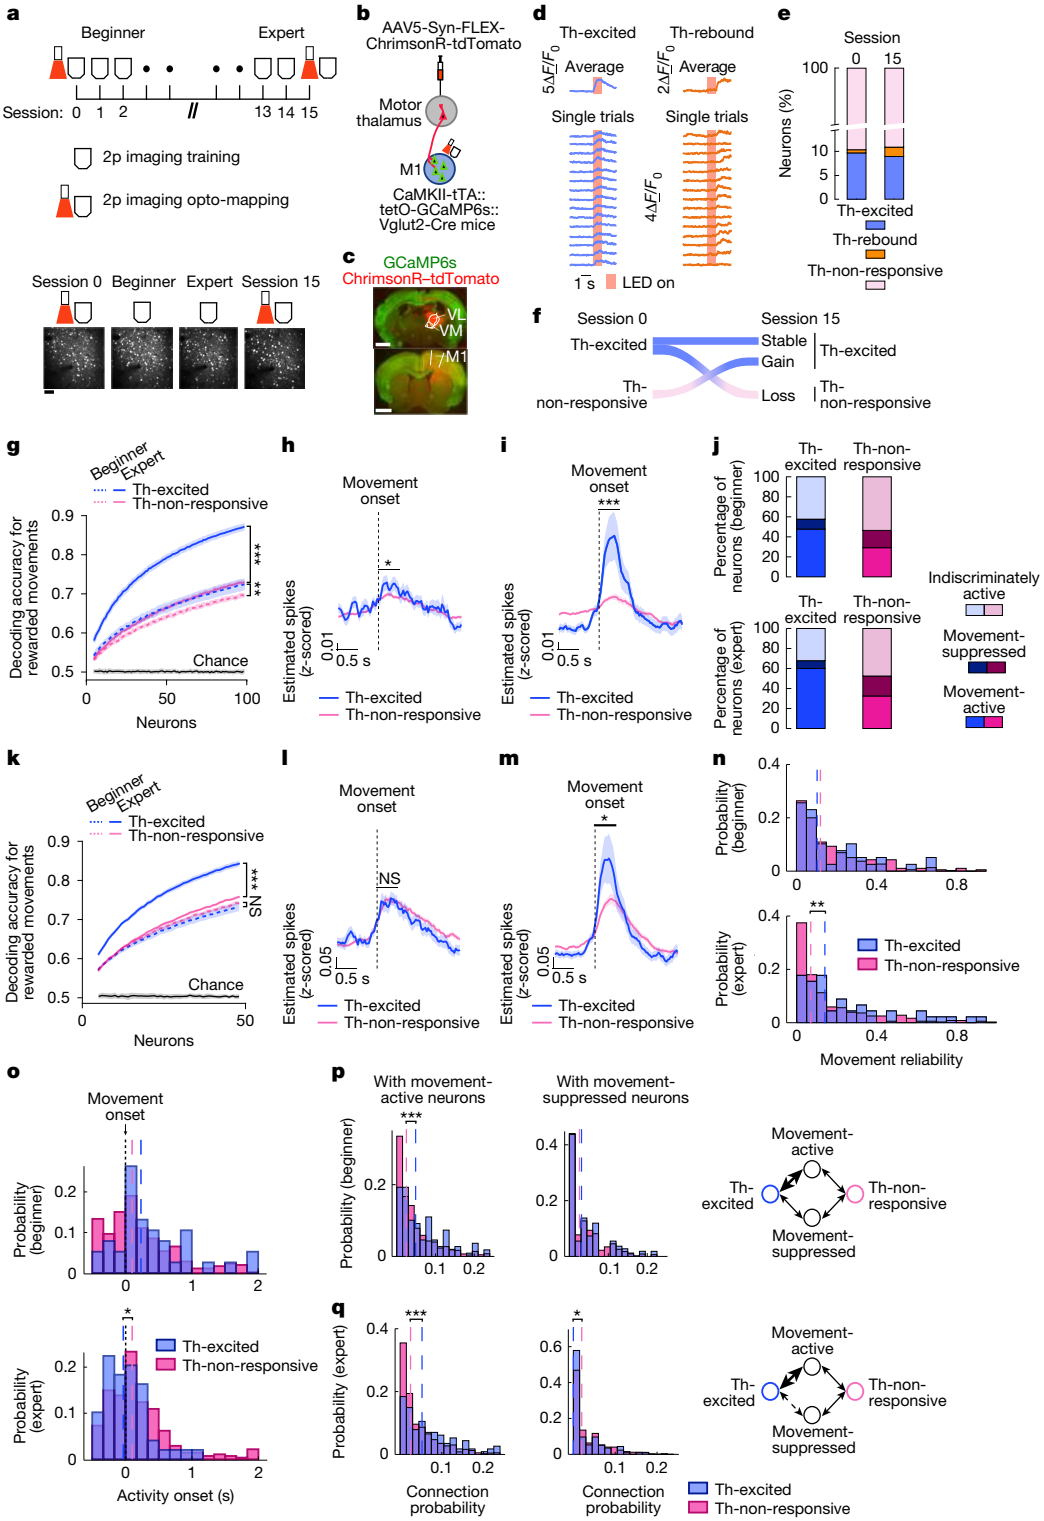

Fig. 2 | See next page for caption.

the activity onset of Th-excited movement-active neurons was not earlier than that of Th-non-responsive movement-active neurons (Fig. 2o, top), but it shifted earlier in experts, such that a majority of Th-excited movement-active neurons in experts started their activity before or around movement onset (Fig. 2o, bottom).

To summarize these results, over learning, Th-excited neurons become more strongly engaged during rewarded movements and shift their activity timing earlier relative to movement onset. Together with the similar observations in the activity of thalamic inputs (Fig. 1), these results support the notion that, with learning, the thalamocortical pathway becomes the key driver of M1 L2/3 activity underlying the learned movement. Specifically, we propose that the thalamic inputs drive a subset of movement-active neurons in M1 L2/3, which will then propagate the activity to other movement-active neurons in M1 through local connectivity.

Therefore, we hypothesized that Th-excited neurons have dense and specific connections to nearby neurons. However, we do not have a direct readout of synaptic connectivity of imaged neurons. Instead,

Fig. 2 | Thalamus-excited M1 L2/3 neurons preferentially encode learned movements in experts. a, Top, experimental timeline. Bottom, max-intensity projection of in vivo two-photon (2p) fluorescence images of the same FOV in M1 L2/3 imaged at the pre-training (session 0) and post-training (session 15) opto-mapping sessions and the first (session 1, beginner) and last (session 14, expert) training sessions. Scale bar, 50 µm. **b**, Experimental strategy. ChrimsonR is expressed in thalamic neurons and imaging and optogenetic stimulation are done in M1. **c**, Coronal sections of the motor thalamus (top) and M1 (bottom) showing the expression of ChrimsonR–tdTomato in thalamic neurons and GCaMP6s in cortical neurons. Scale bars, 1 mm. **d**, Bottom, single-trial responses of example Th-excited (left) and Th-rebound (right) neurons. Top, trial averages. **e**, Fractions of Th-excited, Th-rebound and Th-non-responsive neurons during pre-training (session 0) and post-training (session 15) opto-mapping sessions ($n = 7$ mice; 1,494 and 1,485 neurons for sessions 0 and 15, respectively). **f**, Reorganization of individual neurons' responses to the stimulation of thalamic inputs during pre-training (session 0) and post-training (session 15) opto-mapping sessions. Three categories of Th-excited neurons. About half of Th-excited neurons were Th-excited in both sessions (stable, $n = 47$ neurons), whereas the rest were only Th-excited in pre-training (loss, $n = 51$) or post-training (gain, $n = 49$) sessions. **g**, Cross-validated decoding accuracy for distinguishing rewarded movement and non-movement time points, computed similarly to Fig. 1h. Th-excited neurons outperformed Th-non-responsive neurons in experts (t-test, $P < 10^{-12}$) and in beginners (t-test, $P < 0.01$). Shaded areas show 95% confidence intervals of the mean over cross-validation folds. **h**, Trial-average activity of all Th-excited neurons (blue; 111 neurons) and Th-non-responsive neurons (pink; 1,719 neurons) from beginner sessions aligned to rewarded movement onset (dashed line). Mean ± s.e.m. $P < 0.05$, rank-sum test on baseline (−1.5 s to −1 s before movement onset) subtracted movement-related (0–0.5 s after movement onset) activity. **i**, Same as **h** but for experts (115 Th-excited neurons, 1,365 Th-non-responsive neurons). $P < 0.001$, rank-sum test. **j**, Fractions of movement-active, movement-suppressed and indiscriminately active neurons within the Th-excited and Th-non-responsive populations in beginners (top) and experts (bottom). **k**, Same as **g** but for movement-active neurons. Th-excited movement-active neurons outperformed Th-non-responsive movement-active neurons in experts (t-test, $P < 10^{-9}$) but not in beginners (t-test, $P = 0.128$; NS, not significant). Shaded areas show 95% confidence intervals of the mean over cross-validation folds. **l**, Same as **h** but for movement-active neurons (53 Th-excited neurons, 501 Th-non-responsive neurons). $P = 0.49$, rank-sum test. **m**, Same as **i** but for movement-active neurons (69 Th-excited neurons, 444 Th-non-responsive neurons). $P < 0.05$, rank-sum test. **n**, Histograms of the movement reliability of individual neurons, defined as the fraction of rewarded movements during which the neuron is active. Top, no significant difference in the movement reliability of Th-excited movement-active neurons and Th-non-responsive movement-active neurons in beginners ($P = 0.93$, Wilcoxon signed-rank test; signed-rank test). Bottom, Th-excited movement-active neurons have higher reliability than Th-non-responsive movement-active neurons in experts ($P < 0.01$, signed-rank test). **o**, Histograms of the activity onset of individual movement-active neurons. Top, no significant difference between Th-excited movement-active neurons and Th-non-responsive movement-active neurons in beginners ($P = 0.05$, signed-rank test). Bottom, the activity onset of Th-excited movement-active neurons is earlier than that of Th-non-responsive movementactive neurons and precedes movements in experts ($P < 0.05$, signed-rank test). Vertical broken lines indicate the median. **p**,**q**, Correlation-based functional connectivity. **p**, In beginners, Th-excited neurons have a higher connection probability than Th-non-responsive neurons with movementactive neurons ($P < 0.001$, signed-rank test) but not with movement-suppressed neurons ($P = 0.805$). Right, schematic representing connection probabilities at beginner sessions. **q**, In experts, Th-excited neurons have a higher connection probability than Th-non-responsive neurons with movement-active neurons ($P < 0.001$, signed-rank test) but a lower connection probability with movementsuppressed neurons ($P < 0.05$, signed-rank test). All tests are two-sided. For panels **g**–**i**, **k** and **m**–**q**, $^{*}P < 0.05$, $^{**}P < 0.01$ and $^{***}P < 0.001$.

we inferred their functional connectivity using a correlation-based connectome inference method (FARCI)[21]. We found that Th-excited neurons had an overall higher connection probability to other neurons at both the beginner and the expert stages (Fig. 2p,q and Extended Data Fig. 9). This higher connection probability of Th-excited neurons was specifically to movement-active neurons. By contrast, the connection probability to movement-suppressed neurons was not different between Th-excited neurons and Th-non-responsive neurons in beginners. In expert sessions, the connection probability to movement-suppressed neurons was significantly lower for Th-excited neurons than for Th-non-responsive neurons. Thus, Th-excited neurons are specifically connected to movement-active neurons. In summary, the functional connectivity patterns of Th-excited neurons support their role in selectively relaying thalamic input activity to the movement-active population in M1 L2/3.

We next addressed the question of whether thalamic inputs in experts selectively activate Th-excited neurons, or whether Th-excited neurons are generally more excitable, such that any stimulation tends to activate them. To distinguish between these possibilities, we performed analogous stimulation experiments for the other inputs (S1, cM1 and M2; Extended Data Fig. 10a).

Stimulation of inputs from S1, cM1 and M2 activated substantially fewer neurons in M1 L2/3 than did thalamic input stimulation (Extended Data Fig. 10b). Furthermore, the few S1-, cM1- and M2-excited neurons did not show enhanced movement-related activity in expert sessions, in contrast to Th-excited neurons (Extended Data Fig. 10c–h). These results underscore the unique role of thalamic inputs in driving movement-encoding M1 L2/3 neurons in experts.

Specific encoding in Th-excited neurons

The results above suggest that Th-excited neurons in experts preferentially encode the learned movement. We sought to define the specific behavioural features encoded by the population of Th-excited neurons. However, this was challenging, in part because of the small number of Th-excited neurons in each imaging field. To overcome this issue, we developed shared representation discovery (ShaReD), a computational technique that extends canonical correlation analysis (CCA) for aligning high-dimensional paired multimodal data across individuals (Fig. 3a). In short, ShaReD identifies common feature combinations across all individuals in one modality (behavioural features) that are maximally correlated with specific feature combinations in another modality (neural activity) in each individual. Essentially, ShaReD identifies consistent projections of the behavioural data (using weights for behavioural features) common to all individuals that correlate with projections of the neural data in each individual (using weights assigned to neurons). This allowed us to use the neural activity from all mice to identify the same behavioural features, thus overcoming the limited number of neurons in each mouse.

The behavioural features we considered were the lever position and lever speed at 21 temporal delays between −1 s and +1 s relative to neural activity. Thus, the maximum dimensionality of behavioural features considered was 42. Applying ShaReD to movement-active neurons in the Th-excited and Th-non-responsive populations separately in expert-session data, we identified the behavioural components that were highly correlated with neural activity across all mice. We focused on the first two components with the highest correlations (Extended Data Fig. 11b). The first behavioural component for Th-excited neurons corresponded mostly to a combination of lever position and lever speed at positive lags of around 200 ms, indicating that the activity of Th-excited neurons preceded movements by around 200 ms (Fig. 3b). Conversely, the first behavioural component of the Th-non-responsive neurons was broader in time and peaked at a zero-to-slightly-negative delay, implying that the activity of these neurons relates more to current or immediate past movements (Fig. 3b). The second behavioural component was related mostly to a combination of future lever position and

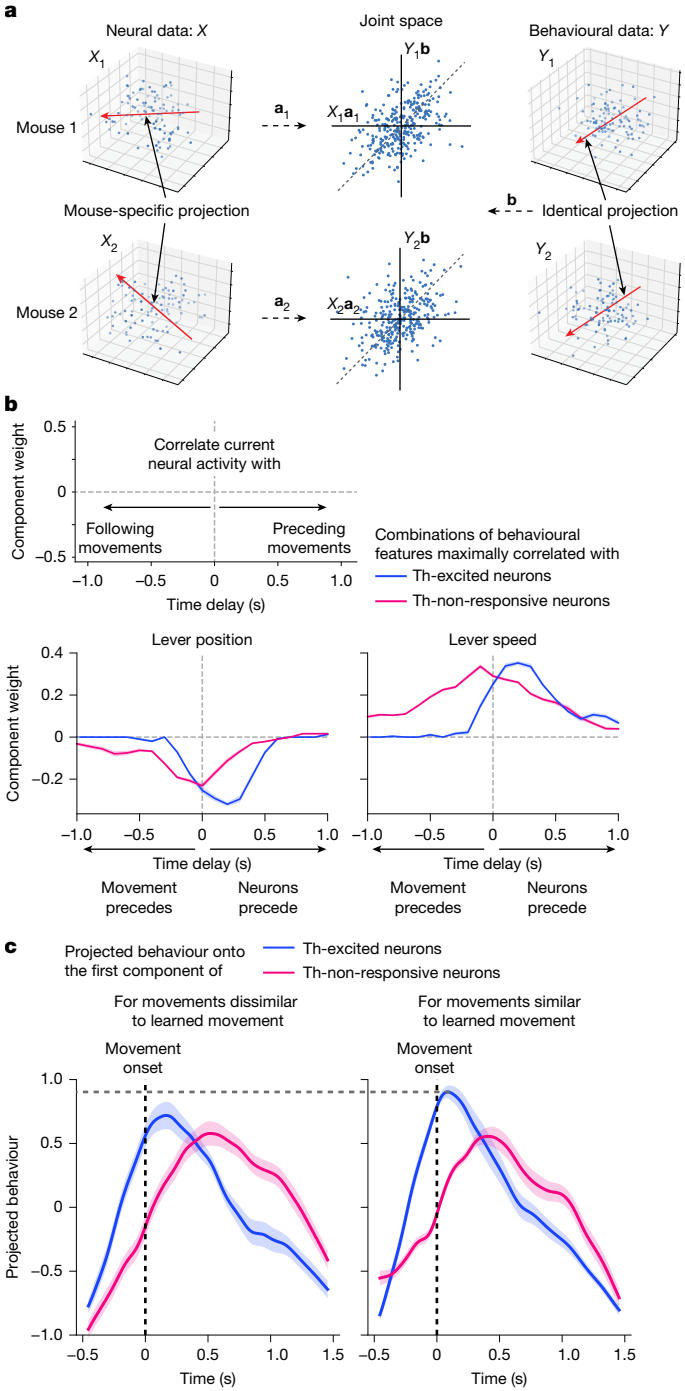

Fig. 3 | ShaReD reveals that Th-excited neurons in experts encode upcoming learned movements. a, Schematic of ShaReD, an extension of CCA for generalizing across multiple animals. Using synthetic datasets for two individuals (i), we visualize neural data (X_i) and behavioural data (Y_i) in three dimensions. Each axis represents one neuron (left column) or one behavioural feature (right column). ShaReD finds the optimal individual neural projection vector $\mathbf{a}_i$ and behavioural projection vector $\mathbf{b}$ that maximize the correlation between $X_i\mathbf{a}_i$ and $Y_i\mathbf{b}$ across all individuals. Red arrows indicate projection vectors for the first component. If the projected data is correlated, the paired points will align near the diagonal in the joint space (middle column). **b**, Weights (unit L2-norm) for the first ShaReD behavioural component for movement-active Th-excited neurons (blue) or Th-non-responsive (pink) neurons. Shaded areas show 95% confidence intervals of the mean over cross-validation folds. For Th-non-responsive neurons, weights are computed from 50 neuron samplings to match the counts of Th-excited neurons and then averaged. The first behavioural component for Th-excited neurons corresponds mainly to a combination of the lever position and the lever speed for movements following the neural activity by around 200 ms, whereas Th-non-responsive neurons correlate more with past movements. Although overall weight signs are arbitrary, the relative signs between position and speed weights must be maintained to preserve their relationship with the underlying movement dynamics. **c**, Projections of non-stereotypical movements (correlation < 0.4 with learned movement pattern; left) and stereotypical movements (correlation > 0.7; right) onto the first behavioural component identified for Th-excited (blue) and Th-non-responsive (pink) neurons. The learned movement pattern was defined as the average lever position trace from movement onset to 1 s after onset, using a randomly selected half of rewarded movements per session. Mean ± s.d. from ten iterations of random sampling. Th-excited neurons preferentially encode upcoming learned movements.

past lever speed for both Th-excited and Th-non-responsive neurons (Extended Data Fig. 11a). To evaluate the specificity of the temporal relationship that we discovered with ShaReD—in which the activity of movement-active Th-excited neurons precedes behavioural features by around 200 ms—we performed a further analysis. We temporally shifted the behavioural component weights (which already capture this approximately 200 ms precedence) and decoded the shifted behavioural projections using the activity of Th-excited neurons (Extended Data Fig. 11c). The decoding accuracy peaked at a shift of zero, confirming that the leading temporal relationship identified by ShaReD was indeed optimal.

We next investigated the type of movements that these behavioural components represent, and how the behavioural components relate to the learned movement. We projected rewarded movements onto the first behavioural components of Th-excited and Th-non-responsive neurons (Fig. 3c). We analysed these projections separately for movements with low and high similarity to the learned movement pattern (Methods). The projections onto the first behavioural component of Th-excited neurons consistently preceded those of Th-non-responsive neurons across all movements. Notably, for movements that were highly similar to the learned movement, the Th-excited projections exhibited a stronger modulation amplitude (Fig. 3c, right), compared with low-similarity movements (Fig. 3c, left). Thus, the combination of behavioural features that best correlates with the activity of Th-excited neurons is maximized near the onset of the learned movement. These observations suggest that Th-excited neurons in experts encode the learned movement, with substantial activity preceding the movement.

Refinement of the thalamic influence on M1

We have shown that Th-excited neurons in experts preferentially encode the learned movement, and that this preferential encoding emerges during learning. The improvement during learning could arise from two scenarios, which could co-exist. In the first scenario, the same neurons are Th-excited throughout learning, but they improve their encoding of the learned movement during learning. In the second scenario, different neurons are Th-excited before and after learning, with the neurons encoding the learned movement becoming Th-excited, and/or the neurons with poor encoding losing their responsiveness for thalamic inputs. We evaluated these scenarios by examining the movement-related activity of stable, loss, and gain Th-excited neurons as defined above (Figs. 2f and 4a). In this nomenclature, stable and loss neurons are Th-excited in beginners, and stable and gain neurons are Th-excited in experts. Discriminating between rewarded movement and non-movement time points, we found that, notably, the population that best encoded rewarded movements in beginners was the gain neurons that were not Th-excited in beginners, whereas the loss neurons had the lowest decoding accuracy (Fig. 4b, left). In experts, the stable and gain neurons substantially improved their encoding

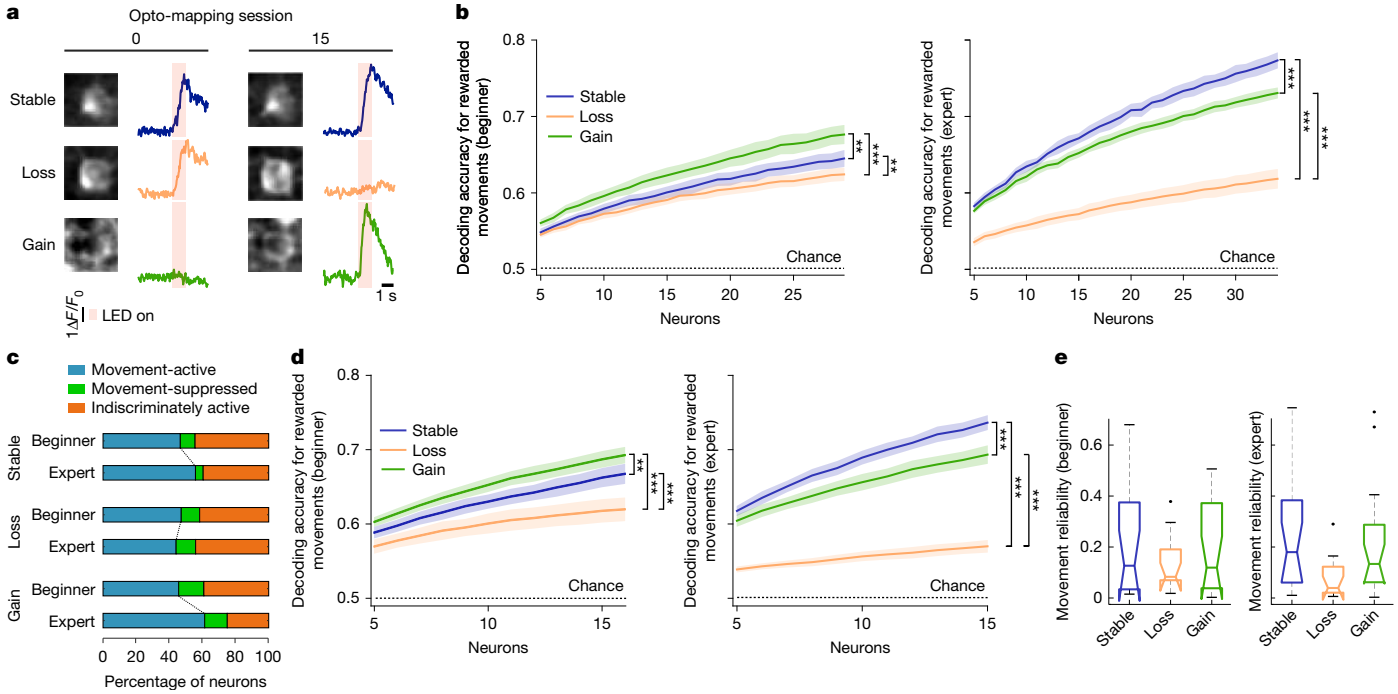

Fig. 4 | The learning-related reorganization of thalamic influence on M1 leads to the encoding of learned movements by Th-excited neurons in experts. a, Example longitudinally recorded neurons and their trial-averaged responses to the optogenetic stimulation of thalamic inputs. One each of stable, loss and gain Th-excited neurons is shown. **b**, Cross-validated decoding accuracy for discriminating between rewarded movement and non-movement time points, similar to Fig. 1h. Shaded areas show 95% confidence intervals of the mean over cross-validation folds. Left, in beginners, gain neurons had a higher decoding accuracy than stable and loss neurons (pairwise *t*-test with Holm–Bonferroni correction, $P < 0.01$ for gain versus stable, $P < 0.01$ for loss versus stable, $P < 10^{-4}$ for gain versus loss). Right, in experts, stable and gain neurons were higher than loss neurons (pairwise *t*-test with Holm–Bonferroni correction, $P < 10^{-11}$ for loss versus stable, $P < 10^{-8}$ for loss versus gain). In addition, stable neurons were higher than gain neurons ($P < 10^{-5}$). **c**, The fractions of stable, loss and gain neurons that were movement-active, movement-suppressed and indiscriminately active in beginner and expert sessions. **d**, Same as **b** but for movement-active neurons. Left, in beginners, both stable and gain movement-active neurons were better than loss movement-active neurons (pairwise *t*-test with Holm–Bonferroni correction, $P < 0.001$ and $P < 10^{-5}$). Gain movement-active neurons were better than stable neurons ($P < 0.01$). Right, in experts, stable and gain movement-active neurons were better than loss movement-active neurons (pairwise *t*-test with Holm–Bonferroni correction, $P < 10^{-13}$ and $P < 10^{-10}$). Stable movement-active neurons were better than gain neurons ($P < 10^{-4}$). **e**, Movement reliability of stable, loss and gain Th-excited movement-active neurons in beginners (left) and experts (right). Box plot elements: centre line, median; box limits, upper and lower quartiles; whiskers, 1.5× interquartile range; points, outliers. For panels **b** and **d**, $^{**}P < 0.01$ and $^{***}P < 0.001$.

of rewarded movements during learning and clearly outperformed the loss neurons (Fig. 4b, right). The encoding of loss neurons did not improve over learning. The fraction of movement-active neurons increased in the stable and gain populations, but not in the loss neurons, partially explaining the decoding results above (Fig. 4c). However, even when we focused our decoding analysis on movement-active neurons of each population, the same trends were observed. Namely, in beginners, decoding by movement-active neurons was best in the gain population, followed by the stable and then the loss neurons (Fig. 4d, left). In experts, the stable movement-active neurons outperformed the gain population, which, in turn, was better than the loss population (Fig. 4d, right). In addition, in experts, stable and gain movement-active neurons had a higher reliability for rewarded movements (Fig. 4e) and an earlier activity onset relative to movement onset than did loss movement-active neurons (Extended Data Fig. 12b), although these differences did not reach statistical significance, probably because of the limited numbers of neurons in each category. These results indicate that the learning-related reorganization of the thalamic influence on M1 involves both of the two scenarios discussed. Namely, the stably Th-excited neurons improve their movement encoding during learning. Furthermore, some of the neurons that are strongly movement-encoding but not Th-excited in beginners become Th-excited in experts. Conversely, some of the neurons with poor movement encoding lose their responsiveness to thalamic inputs as learning progresses.

Th inputs are required in experts but not in beginners

The results so far suggest that a specific drive of M1 movement-active neurons by thalamic inputs is crucial for generating learned movements. To assess the functional importance of thalamic activity in the expert behaviour, we examined the effect of inactivating the motor thalamus on behaviour.

To inactivate motor-thalamus neurons, we unilaterally expressed inhibitory DREADDs (hM4Di[22]) or tdTomato as a control in thalamic neurons by injecting AAV-hSyn-DIO-hM4D-mCherry (or AAV-FLEX-tdTomato) into the motor thalamus of Vglut2-Cre mice (Fig. 5a). These mice were trained for two sessions without manipulation. Then, on day 3, we acutely inactivated the motor thalamus by intraperitoneal (i.p.) injection of clozapine-*N*-oxide (CNO) (Fig. 5b). This inactivation had no significant effect on performance in these beginner mice (Fig. 5c). By contrast, when we trained another group of mice for 13 sessions to reach the expert stage and then inactivated the motor thalamus by i.p. injection of CNO, their performance was significantly impaired, as indicated by fewer rewarded trials and increased time from cue to reward in the hM4Di group (Fig. 5d,e). Thus, motor-thalamus activity is crucial for the expert but not for the beginner behaviour. This aligns with the refinement of thalamic influence on M1 during learning described above.

The effects of motor-thalamus inactivation could in theory be mediated by thalamic projections to brain areas other than M1. To directly

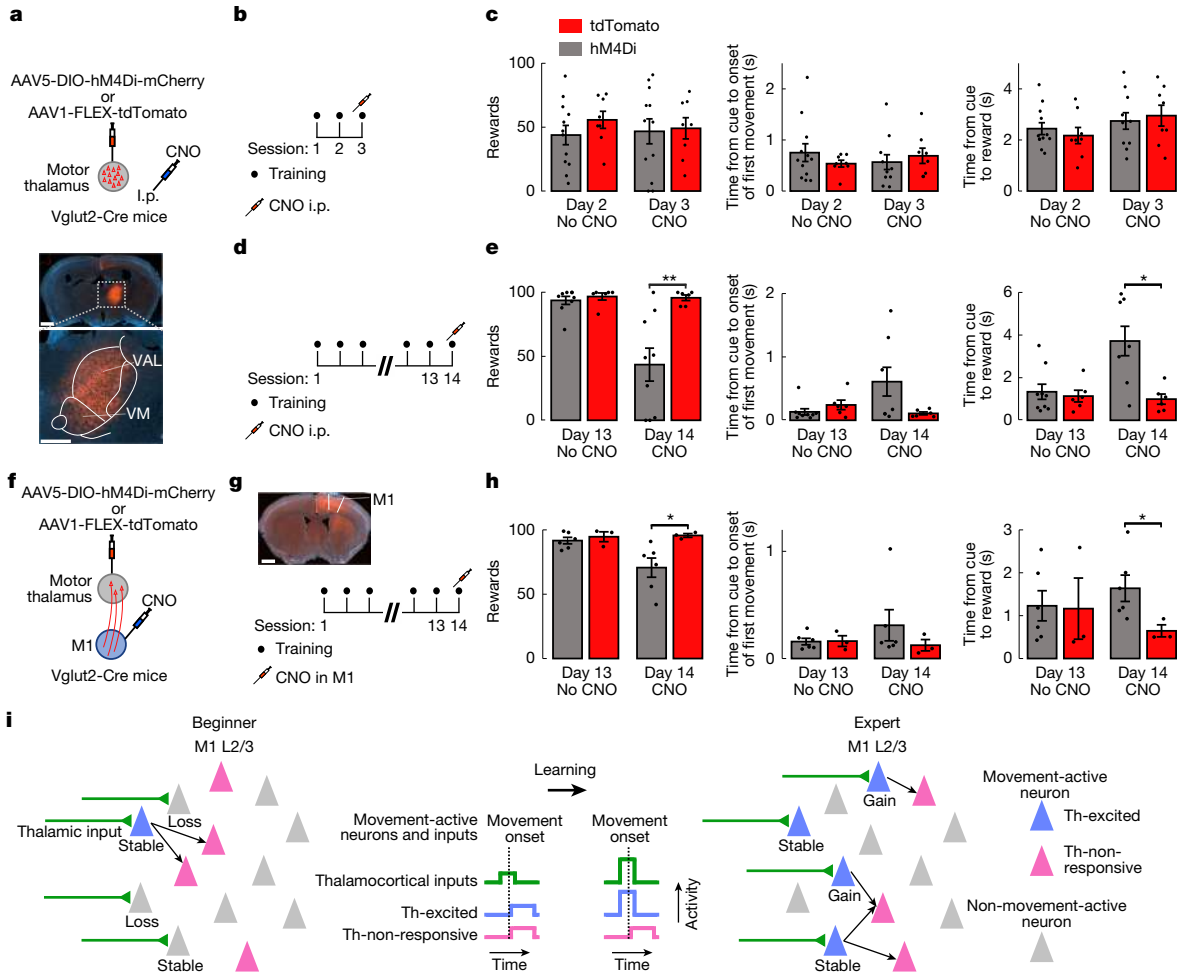

Fig. 5 | Thalamic activity is required for the execution of learned movements. a, Top, experimental strategy. hM4Di-mCherry or tdTomato was expressed in motor thalamic neurons and CNO was administered i.p. Bottom, coronal sections showing hM4Di-mCherry expression in the motor thalamus. Blue, DAPI. Scale bars, 1 mm (top); 0.5 mm (bottom). **b**, Experimental timeline for motor-thalamus inactivation in beginners. **c**, No significant difference in the number of rewarded trials (left; $P = 1$, rank-sum test), the time from cue to reward (right; $P = 0.74$, rank-sum test) and the time from cue onset to the onset of the first movement (middle; $P = 0.35$, rank-sum test) in beginner mice. $n = 12$ hM4Di mice and $n = 8$ tdTomato mice. Circles correspond to individual mice. **d**, Experimental timeline for motor-thalamus inactivation in experts. **e**, CNO administration in experts resulted in a significantly lower number of rewarded trials (left, $P < 0.01$, rank-sum test), a significantly longer time between cue and reward (right, $P < 0.05$, rank-sum test) and a non-significant increase in the time from cue onset to movement onset (middle, $P = 0.23$, rank-sum test) in the hM4Di group ($n = 9$ mice) compared with the tdTomato group ($n = 6$ mice). Circles correspond to individual mice. **f**, Experimental strategy. hM4Di-mCherry or tdTomato was expressed in motor thalamic neurons and CNO was administered locally to M1. **g**, Top, coronal sections showing hM4Di-mCherry expression in thalamic axons in M1. Scale bar, 1 mm. Bottom, experimental timeline for inactivation of thalamic axons in M1 in experts. **h**, Local administration of CNO in M1 in experts resulted in a significantly lower number of rewarded trials (left; $P < 0.05$, rank-sum test), a significantly longer time between cue and reward (right; $P < 0.05$, rank-sum test), and a non-significant increase in the time from cue onset to movement onset (middle; $P = 0.26$, rank-sum test) in the hM4Di group ($n = 6$ mice) compared with the tdTomato group ($n = 3$ mice). Circles correspond to individual mice. All tests are two-sided. **i**, Model schematic. Arrows represent functionally connected neurons. For panels **e** and **h**, $^*P < 0.05$ and $^{**}P < 0.01$.

examine the effect of motor-thalamus inputs to M1, we selectively inactivated motor-thalamus axons in M1. Using the same approach, we expressed hM4Di or tdTomato in the motor thalamus (Fig. 5f), and trained these mice for 13 days to expert level. We then locally administered CNO into M1 on session 14 (Fig. 5g). This manipulation, similarly, resulted in fewer rewarded trials and increased cue-to-reward times in the hM4Di group (Fig. 5h). Thus, motor-thalamus activity—and specifically, its inputs to M1—is necessary for the precise execution of learned movements in experts.

Discussion

Our results reveal that thalamic inputs to M1 are an essential driver for the initiation and execution of learned movements. Our findings are consistent with previous studies highlighting the importance of thalamic inputs in controlling voluntary movements[23–27]. Furthermore, our longitudinal analysis during learning provides a new perspective on how this learned circuit emerges from a precise change in the cellular-level influence of the motor thalamus on M1.

Thalamic inputs to M1 in learned movements

Here we provide several lines of evidence to support a model in which the brain generates the M1 activity pattern to execute learned movements. In this model, immediately before the initiation of the movement, the motor thalamus activates a small group of core neurons in M1. This specific excitation gets amplified and transformed by selective connectivity in M1, leading to the spatio-temporal population activity that drives the learned movement.

The degree of autonomy of M1 in producing the movement-generating activity pattern in M1 has been debated[23,28–31]. In one extreme, M1 could function as an autonomous dynamical system, which generates complex activity patterns through M1-intrinsic connectivity, perhaps in response to non-specific inputs. In the other extreme, M1 could be a purely passive machine that simply reflects the specific and time-varying input without transforming it. Our results do not support either of these extremes. Rather, we propose that specific and time-varying inputs from the motor thalamus interact with local connectivity in M1 to produce the movement-generating activity pattern. The balance between inputs and intrinsic connectivity might depend on the complexity of the learned skill: longer or more intricate sequences might involve multiple local circuits, each driven by specific thalamic inputs, forming a sequence of motor modules that collectively generate the skill.

Our results extend previous studies highlighting the role of thalamic inputs in enabling M1 to initiate learned movements[32,33]. Our approach uncovered a small fraction (around 9%) of M1 L2/3 neurons that are strongly driven by thalamocortical inputs. Although these Th-excited neurons might not all be monosynaptically driven by thalamic inputs, they are still likely to be excited by local effects of thalamic inputs in M1, because Th-excited neurons were not excited by stimulation of the other main inputs (Extended Data Fig. 10). It is also likely that our approach only reveals strong connections, and excitation that is below the threshold for spike generation or below the limit of detection by calcium imaging would remain undetected. Furthermore, thalamic inputs to M1 L2/3 come mostly from the motor thalamic nuclei, including VAL and VM, and from nearby nuclei such as VPM, VPL and PO (Extended Data Fig. 3). Dissociating potentially distinct functions of M1 inputs from these nuclei would provide more detailed descriptions of the neural circuits that underlie learned movements[24,25,34]. Regardless, our model resembles what was proposed in a study in zebra finches[35], which suggested that the motor thalamic nucleus uvaeformis activates a subset of neurons in the zebra finch song nucleus HVC (homologous to the primary or premotor motor cortex in mammals[36]) to initiate singing. The conservation of this circuit motif across evolution underscores its importance for learned behaviours.

Motor learning sculpts the thalamus–M1 circuit

Our longitudinal investigation identified a number of learning-related changes that contribute to the operation of the expert circuit. First, the amount of movement-related information in the thalamic input activity increased. Second, possibly as a result of the first, the stably Th-excited neurons increased their movement-related activity. We note that plasticity in downstream circuits could also contribute to this change. Third, some of the neurons with strong movement-related activity in beginners that were not Th-excited gained Th responsiveness. Fourth, some of the Th-excited neurons in beginners with low movement-related activity lost their Th responsiveness after learning. Finally, the functional connectivity within M1 became more refined. These parallel changes ensure that thalamic inputs selectively target the correct neurons to drive the learned movement.

The third and fourth changes suggest that the motor thalamus learns to activate the right neurons and to stop driving the wrong neurons. Such a reorganization probably requires precise plasticity during learning. Indeed, motor learning induces plasticity of both thalamocortical axonal boutons innervating M1 (ref. 32) and local M1 dendritic spines[1,4,24,37,38], and spine plasticity of M1 L2/3 neurons during motor learning has been shown to be directed to form synaptic clusters dedicated to the learned movement[4,5]. These synaptic changes coincide with large-scale changes in the transcriptional states of M1 L2/3 neurons[39]. We propose that specific, targeted thalamocortical plasticity refines the thalamocortical pathway so that the motor thalamus activates the appropriate subset of M1 L2/3 neurons to initiate the process of generating learned movements. Other mechanisms, such as cell-intrinsic changes in excitability[40], might also contribute to such learning-related reorganizations. The parallel refinement of the thalamocortical pathway and local recurrent connectivity within M1 results in a precise circuit to generate a reproducible pattern of activity dedicated to the learned movement (Fig. 5i), composed of Th-excited neurons that are functionally connected to movement-active neurons. This functional connectivity might be dynamically gated by network states, which could explain why our stimulation during the opto-mapping sessions did not lead to full-blown movement-generating population activity.

Notably, inactivation of the motor thalamus did not affect this behaviour in beginner mice. It is possible that a more complete inactivation could lead to behavioural effects in beginners, but the same manipulation significantly impaired the behaviour in experts, supporting an increasing role of the motor thalamus in movement generation as learning progresses. We also note that previous studies have found that extensive training over months can make the movements that are at first M1-dependent become M1-independent[10,11]. The expert stage in our current study is a stage at which M1 is still crucial for the behaviour. These observations highlight that the brain has multiple circuits that can generate similar behaviours, and the degree of learning is a determinant of which circuit is used. Furthermore, other long-range inputs that we examined also exhibited learning-related changes in activity. Future studies should dissect the specific contributions of these inputs to the dynamics and plasticity of M1 circuits during motor learning.

ShaReD

To understand the relationship between the ensemble activity of the small number of Th-excited neurons per animal and behaviour, we developed an analytical approach that we termed ShaReD. This technique builds on CCA[41], which identifies correlations between neural features and behavioural features, allowing both the behavioural feature space and the space of all neurons' activities to be high-dimensional. This high dimensionality calls for a systematic investigation of the relationship between neural population activity and combinations of behavioural features. Although previous methods have examined these relationships by jointly modelling neural and behavioural dynamics, they often lack generalizability across individuals[42,43], or necessitate identical behaviour between individuals[44]. ShaReD takes a different approach by identifying a single behavioural projection shared across individuals.

Applying ShaReD to our data revealed specific behavioural subspaces (that is, combinations of lever position and speed at different time lags) that were commonly correlated with individual subspaces of Th-excited neuron activity across individuals. These specific subspaces within the high-dimensional behavioural space would have been difficult to identify without the unsupervised approach facilitated by ShaReD. We present ShaReD as a generalizable dimensionality reduction technique to identify specific behavioural modes that are shared across individuals and to learn how they relate to patterns of neural activity in different individuals.

ShaReD operates on the assumption that there are common combinations of behavioural features encoded across individuals, and identifies these common components. Thus, it is designed to mostly ignore variability of behavioural encoding across individuals. This is a useful property, but also a limitation that should be considered, and we stress the importance of using multiple analytical approaches to investigate how neural activities represent behaviours that might be heterogeneous across individuals.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41586-025-08962-8.

1. Peters, A. J., Chen, S. X. & Komiyama, T. Emergence of reproducible spatiotemporal activity during motor learning. *Nature* **510**, 263–267 (2014).
2. Peters, A. J., Liu, H. & Komiyama, T. Learning in the rodent motor cortex. *Annu. Rev. Neurosci.* **40**, 77–97 (2017).
3. Guo, L. et al. Dynamic rewiring of neural circuits in the motor cortex in mouse models of Parkinson's disease. *Nat. Neurosci.* **18**, 1299–1309 (2015).
4. Hedrick, N. G. et al. Learning binds new inputs into functional synaptic clusters via spinogenesis. *Nat. Neurosci.* **25**, 726–737 (2022).
5. Hedrick, N. G., Wright, W. J. & Komiyama, T. Local and global predictors of synapse elimination during motor learning. *Sci. Adv.* **10**, eadk0540 (2024).
6. Chen, S. X., Kim, A. N., Peters, A. J. & Komiyama, T. Subtype-specific plasticity of inhibitory circuits in motor cortex during motor learning. *Nat. Neurosci.* **18**, 1109–1115 (2015).
7. Makino, H., Hwang, E. J., Hedrick, N. G. & Komiyama, T. Circuit mechanisms of sensorimotor learning. *Neuron* **92**, 705–721 (2016).
8. Sanes, J. N. & Donoghue, J. P. Plasticity and primary motor cortex. *Annu. Rev. Neurosci.* **23**, 393–415 (2000).
9. Kawai, R. et al. Motor cortex is required for learning but not for executing a motor skill. *Neuron* **86**, 800–812 (2015).
10. Hwang, E. J. et al. Disengagement of motor cortex from movement control during long-term learning. *Sci. Adv.* **5**, eaay0001 (2019).
11. Hwang, E. J., Dahlen, J. E., Mukundan, M. & Komiyama, T. Disengagement of motor cortex during long-term learning tracks the performance level of learned movements. *J. Neurosci.* **41**, 7029–7047 (2021).
12. Weiler, N., Wood, L., Yu, J., Solla, S. A. & Shepherd, G. M. G. Top-down laminar organization of the excitatory network in motor cortex. *Nat. Neurosci.* **11**, 360–366 (2008).
13. Itokazu, T. et al. Streamlined sensory motor communication through cortical reciprocal connectivity in a visually guided eye movement task. *Nat. Commun.* **9**, 338 (2018).
14. Hooks, B. M. et al. Organization of cortical and thalamic input to pyramidal neurons in mouse motor cortex. *J. Neurosci.* **33**, 748–760 (2013).
15. Mao, T. et al. Long-range neuronal circuits underlying the interaction between sensory and motor cortex. *Neuron* **72**, 111–123 (2011).
16. Muñoz-Castañeda, R. et al. Cellular anatomy of the mouse primary motor cortex. *Nature* **598**, 159–166 (2021).
17. Broussard, G. J. et al. In vivo measurement of afferent activity with axon-specific calcium imaging. *Nat. Neurosci.* **21**, 1272–1280 (2018).
18. Klapoetke, N. C. et al. Independent optical excitation of distinct neural populations. *Nat. Methods* **11**, 338–346 (2014).
19. Chen, T.-W. et al. Ultra-sensitive fluorescent proteins for imaging neuronal activity. *Nature* **499**, 295–300 (2013).
20. Vong, L. et al. Leptin action on GABAergic neurons prevents obesity and reduces inhibitory tone to POMC neurons. *Neuron* **71**, 142–154 (2011).
21. Meamardoost, S. et al. FARCI: fast and robust connectome inference. *Brain Sci.* **11**, 1556 (2021).
22. Krashes, M. J. et al. Rapid, reversible activation of AgRP neurons drives feeding behavior in mice. *J. Clin. Invest.* **121**, 1424–1428 (2011).
23. Sauerbrei, B. A. et al. Cortical pattern generation during dexterous movement is input-driven. *Nature* **577**, 386–391 (2020).
24. Tanaka, Y. H. et al. Thalamocortical axonal activity in motor cortex exhibits layer-specific dynamics during motor learning. *Neuron* **100**, 244–258 (2018).
25. Dacre, J. et al. A cerebellar–thalamocortical pathway drives behavioral context-dependent movement initiation. *Neuron* **109**, 2326–2338 (2021).
26. Inagaki, H. K. et al. A midbrain–thalamus–cortex circuit reorganizes cortical dynamics to initiate movement. *Cell* **185**, 1065–1081 (2022).
27. Yang, W., Tipparaju, S. L., Chen, G. & Li, N. Thalamus-driven functional populations in frontal cortex support decision-making. *Nat. Neurosci.* **25**, 1339–1352 (2022).
28. Russo, A. A. et al. Motor cortex embeds muscle-like commands in an untangled population response. *Neuron* **97**, 953–966 (2018).
29. Churchland, M. et al. Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).
30. Schimel, M., Kao, T.-C. & Hennequin, G. When and why does motor preparation arise in recurrent neural network models of motor control? *eLife* **12**, RP89131 (2023).
31. Kalidindi, H. T. et al. Rotational dynamics in motor cortex are consistent with a feedback controller. *eLife* **10**, e67256 (2021).
32. Sohn, J. et al. Presynaptic supervision of cortical spine dynamics in motor learning. *Sci. Adv.* **8**, eabm0531 (2022).
33. Takahashi, N. et al. Thalamic input to motor cortex facilitates goal-directed action initiation. *Curr. Biol.* **31**, 4148–4155 (2021).
34. Guo, K., Yamawaki, N., Svoboda, K. & Shepherd, G. M. G. Anterolateral motor cortex connects with a medial subdivision of ventromedial thalamus through cell type-specific circuits, forming an excitatory thalamo-cortico-thalamic loop via layer 1 apical tuft dendrites of layer 5B pyramidal tract type neurons. *J. Neurosci.* **38**, 8787–8797 (2018).
35. Moll, F. W. et al. Thalamus drives vocal onsets in the zebra finch courtship song. *Nature* **616**, 132–136 (2023).
36. Mooney, R. The neurobiology of innate and learned vocalizations in rodents and songbirds. *Curr. Opin. Neurobiol.* **64**, 24–31 (2020).
37. Biane, J. S., Takashima, Y., Scanziani, M., Conner, J. M. & Tuszynski, M. H. Thalamocortical projections exhibit plasticity onto behaviorally-relevant neurons during adult motor learning. *Neuron* **89**, 1173–1179 (2016).
38. Hasegawa, R., Ebina, T., Tanaka, Y. R., Kobayashi, K. & Matsuzaki, M. Structural dynamics and stability of corticocortical and thalamocortical axon terminals during motor learning. *PLoS One* **15**, e0234930 (2020).
39. Dunton, K. L. et al. Divergent learning-related transcriptional states of cortical glutamatergic neurons. *J. Neurosci.* **44**, e0302232023 (2024).
40. Alejandre-García, T., Kim, S., Pérez-Ortega, J. & Yuste, R. Intrinsic excitability mechanisms of neuronal ensemble formation. *eLife* **11**, e77470 (2022).
41. Hotelling, H. in *Breakthroughs in Statistics: Methodology and Distribution* (eds Kotz, S. & Johnson, N. L.) 162–190 (Springer, 1992).
42. Perkins, S. M., Cunningham, J. P., Wang, Q. & Churchland, M. M. Simple decoding of behavior from a complicated neural manifold. *eLife* **12**, RP89421 (2023).
43. Sani, O. G., Abbaspourazad, H., Wong, Y. T., Pesaran, B. & Shanechi, M. M. Modeling behaviorally relevant neural dynamics enabled by preferential subspace identification. *Nat. Neurosci.* **24**, 140–149 (2021).
44. Safaie, M. et al. Preserved neural dynamics across animals performing similar behaviour. *Nature* **623**, 765–771 (2023).

Methods

Mice

All animal procedures were performed in accordance with guidelines and protocols approved by the UCSD Institutional Animal Care and Use Committee and the National Institutes of Health (NIH). C57BL/6 (Carles River), Vglut2-Cre (Vglut2-ires-cre, JAX stock 016963), CaMKII-tTA (B6;CBA-Tg(Camk2a-tTA)1Mmay/J, JAX stock 003010) and tetO-GCaMP6s (B6;DBA-Tg(tetO-GCaMP6s)2Niell/J, JAX stock 024742) mice were group-housed in disposable cages with standard bedding in a temperature-controlled room (around 21 °C) with a reversed light cycle (10:00–22:00: dark). All experiments were performed during the dark cycle. Male and female mice were used at random for experiments.

Surgeries and virus injections

Adult mice (seven weeks or older) were anaesthetized with 1–2% isoflurane and injected with Baytril (10 mg kg^{-1}), dexamethasone (2 mg kg^{-1}) and buprenorphine (0.1 mg kg^{-1}) subcutaneously at the beginning of surgery to prevent infection, inflammation and discomfort. A custom-built head-plate was glued and cemented to the skull. For two-photon imaging experiments, a craniotomy (around 3 mm in diameter) was performed over the right caudal forelimb area of M1, centred at 0.3 mm anterior and 1.5 mm lateral from the bregma. A glass window was implanted over the craniotomy. The edges between the window and the skull were filled with Vetbond (3M). The window was further secured with cyanoacrylate glue and dental acrylic.

For imaging thalamic projections to M1, virus solutions of AAV5-hSyn-FLEX-axon-GCaMP6s (Addgene viral preparation 112010) were injected with glass pipettes into the following coordinates of Vglut2-Cre mice: 1 mm posterior, 1 mm lateral to bregma at a depth of around 3.7 mm from the pia (approximately 0.3 µl). Pipettes were left in the brain for 15 min after each injection to minimize backflow. Imaging was performed around four weeks after surgery.

For imaging S1, cM1 and M2 projections to M1, virus solutions of AAV5-hSyn-FLEX-axon-GCaMP6s and AAV1-CMV-PI-Cre were injected with glass pipettes into the following coordinates: S1: 0.8 mm posterior, 2.1 mm lateral to bregma at a depth of around 0.25–0.5 mm from the pia (approximately 0.1 µl); cM1: 0.3 mm anterior, 1.5 mm lateral to bregma at a depth of around 0.25–0.5 mm from the pia (approximately 0.4 µl); M2 (rostral forelimb area): 2.5 mm anterior, 0.9 mm lateral to bregma at a depth of around 0.25–0.5 mm from the pia (approximately 0.2 µl). Pipettes were left in the brain for 15 min after each injection to avoid backflow. Imaging was performed around four weeks after surgery.

For in vivo two-photon imaging of ChrimsonR axons, virus solutions of AAV1-hSyn-FLEX- GCaMP6f (Addgene viral preparation 100833) and AAV5-hSyn-FLEX-ChrimsonR-tdTomato (Addgene viral preparation 62723) were injected into the following coordinates of Vglut2-Cre mice: 1 mm posterior, 1 mm lateral to bregma at a depth of around 3.7 mm from the pia (approximately 0.4 µl).

For in vivo two-photon imaging of M1 L2/3 neurons combined with optogenetic stimulation of thalamocortical axons, virus solutions of AAV5-hSyn-FLEX-ChrimsonR-tdTomato were injected into the following coordinates of CaMKII-tTA::tetO-GCaMP6s::Vglut2-Cre triple transgenic mice: 1 mm posterior, 1 mm lateral to bregma at a depth of around 3.7 mm from the pia (approximately 0.4 µl). For in vivo two-photon imaging of M1 L2/3 neurons combined with optogenetic stimulation of S1, cM1 and M2 axons, virus solutions of AAV1-hSyn-ChrimsonR-tdTomato were injected into the following coordinates of CaMKII-tTA::tetO-GCaMP6s transgenic mice: S1: 0.8 mm posterior, 2.1 mm lateral to bregma at a depth of around 0.25–0.5 mm from the pia (approximately 0.1 µl); cM1: 0.3 mm anterior, 1.5 mm lateral to bregma at a depth of around 0.25–0.5 mm from the pia (approximately 0.4 µl); M2 (rostral forelimb area): 2.5 mm anterior, 0.9 mm lateral to bregma at a depth of around 0.25–0.5 mm from the pia (approximately 0.2 µl). Experiments were performed around six weeks after surgery.

For motor-thalamus inactivation, virus solutions of AAV5-hSyn-DIO-hM4Di-mCherry (Addgene viral preparation 44362) (diluted 1:5) or AAV1-FLEX-tdTomato (Addgene viral preparation 51503) (control group) were injected into the following coordinates of Vglut2-Cre mice: 1 mm posterior, 1 mm lateral to bregma at a depth of around 3.7 mm from the pia (approximately 0.4 µl). Experiments were performed around four and eight weeks after surgery for i.p. CNO administration and local M1 CNO injection, respectively.

Behaviour

Water restriction started 1–1.5 weeks before the behavioural training. Mice received progressively lower amounts of drinking water until stabilizing at 1 ml per day. Weight was constantly monitored to ensure no more than 30% of the starting body weight was lost. Mice were trained to perform the lever-press task one session per day for 14 days under a microscope. The hardware and software used for behavioural training have been previously described[45]. In brief, the lever comprised a piezoelectric flexible force transducer (LCL-113G, Omega Engineering) attached to a brass rod. The lever position was continuously recorded using LabJack and Ephus (controlled by MATLAB), which worked with custom software running on LabVIEW (National Instruments) that monitored threshold crossing. Dispatcher software controlled the behavioural set-up (Z. Mainen and C. Brody, controlled by MATLAB). Mice used their left paw to grasp a lever attached to a force transducer while resting their right paw on a stationary block. A 6-kHz tone marked a cue period (up to 10 s), during which a successful lever press was rewarded with water (around 10 µl per trial) paired with a 500-ms, 12-kHz tone, and followed by an inter-trial interval (ITI, variable duration of 8-12 s). A successful lever-press movement was defined as crossing two thresholds (around 1.5 mm and 3 mm below the resting position) within 200 ms. Failure to press the lever passing the two thresholds during the cue period triggered a loud white-noise sound and the start of an ITI. Lever presses during ITIs were neither rewarded nor punished. Each session consisted of 100 trials.

Lever-press movement analysis

Movement analyses were performed as previously described[1,4]. In brief, lever displacement traces (voltage recordings from the force transducer) were downsampled from 10 kHz to 1 kHz, then filtered using a four-pole 10 Hz low-pass Butterworth filter, after which the velocity of the lever was determined by smoothing the difference of consecutive points with a moving average window of 5 ms. The envelope of the lever velocity was then extracted using a Hilbert transform, and movement bouts were defined by the envelope crossing a threshold of 4.9 mm per second. Each movement bout was extended by 75 ms on either side. Bouts separated by less than 500 ms were considered continuous. Movement start and end times were defined as the points at which the lever exceeded or fell below the thresholds defined by rest periods before and after the movement bouts. Thresholds were defined as the resting position plus the 99th percentile of the noise distribution, in turn defined as the difference between the Butterworth smoothed trace and the original trace. For reaction-time analysis, the trials in which mice moved the lever within 100 ms before the cue start time were excluded from analysis. Movement correlations within sessions were calculated using the median of all pairwise correlations of rewarded movements that started after cue onset within a single session. Movement correlations across sessions were found using the median correlations of all possible pairs of movements between sessions.

In vivo two-photon imaging

Imaging was performed using a commercial two-photon microscope (B-Scope, ThorLabs) equipped with a ×16/0.8-NA objective (Nikon) and a Ti-Sa laser (Mai Tai) tuned to 925 nm. Image acquisition was controlled

with ScanImage software. Images (512 × 512 pixels at ×2 zoom for neuronal recording or ×5 zoom for axonal recording) were recorded at around 30 Hz for the duration of the behavioural session. Frame times were recorded and synchronized with behavioural recordings using the Ephus software. For some mice, two FOVs were recorded on interleaved days. The FOVs did not overlap.

For axon imaging, we imaged at the most superficial depth that contained many labelled axons. This was almost always within 200 μm of the pia surface. Soma imaging was typically done at a depth of 200–300 μm from the pia surface.

Image analysis

ROI identification and fluorescence analysis. Suite2P software[46] was used to generate regions of interest (ROIs) corresponding to individual neurons or axonal boutons and extract their fluorescence. ROI classifications by the automatic classifier were further refined by manual inspection. Signals from ROIs were deconvolved with a non-negative deconvolution algorithm to remove fluorescence decay and estimate underlying spiking activity[47]. The estimated spikes were used for all neural activity analyses, besides analysing neuronal activity during opto-mapping sessions. For opto-mapping analysis, fluorescence time series were produced by averaging the pixels within each ROI (minus surrounding neuropil) for all imaging frames. The time-varying baseline of a fluorescence trace (F_0) was estimated by smoothing inactive portions of the trace using a previously described iterative procedure[1]. In brief, this process identified the trace's active and inactive portions, removing active portions and using the LOESS-smoothed inactive portions (interpolated across active periods) to estimate the time-varying baseline. The normalized $\Delta F/F_0$ trace was then calculated, in which ΔF was found by subtracting the baseline trace from the raw trace, and F_0 is the calculated time-varying baseline.

To match neurons across multiple imaging sessions, we used the MATLAB code ROIMatchpub (https://github.com/ransona/ROIMatchPub), followed by manual confirmation and corrections after automated detection.

Classification of movement-modulated neurons. Movement-modulated neurons were classified as previously described[1]. In brief, the dot product of the binarized lever trace (movements versus non-movements) and estimated spikes was calculated for each ROI. This value was compared to the dot products when shuffling the movement periods 1,000 times. Actual values above the 97.5 percentile of the shuffled distribution were classified as movement-active, and actual values below the 2.5 percentile were classified as movement-suppressed. All other cells were considered indiscriminately active.

Activity onset. To estimate the timing of the activity onset of individual neurons during rewarded movements, the activity of each neuron (estimated spikes) was aligned to the onsets of rewarded movements. We only considered movements that did not have another movement within 2 s before movement onset to avoid contamination of activity related to previous movements. The activity of each neuron was averaged across all rewarded movements, and the activity onset time was defined by the first frame in which activity was higher than the averaged baseline activity (defined as −2 to −1.5 s before movement onset) by more than one standard deviation and remained above this criterion for 90 ms (three successive frames).

Population activity correlation. The stability of the population across trials was assessed by the correlation of population vectors. Each population vector includes the concatenated activity of all neurons during rewarded movement epochs (2-s time window starting at 0.5 s before the onset of the rewarded movement). The movement–movement activity Pearson's correlation was then calculated between all pairs of all population vectors per FOV.

Correlation-based functional connectivity. Functional connectivity was inferred using the fast and robust connectome inference (FARCI) method[21]. FARCI was designed to infer the functional neuronal connectome from in vivo two-photon Ca imaging. In brief, after steps of estimated spike thresholding (three standard deviations) and smoothing (three frames), FARCI uses partial correlations between pairs of neurons to infer their connectivity. We used the following parameters for spike thresholding, smoothing and partial correlations: 3, 3 and 2, respectively.

Optogenetic mapping

Optogenetic stimulation was performed by delivering 660-nm LED light (Thorlabs, M660L3), with a full width at half maximum of 10 nm, through a bandpass filter (660/10 nm, Thorlabs, FB660-10) and was combined with the two-photon excitation light for imaging with a 750-nm long-pass filter (Chroma T750lpxrxt-UF1). These two lights were directed to the sample by a primary dichroic mirror reflecting 640–700 nm and 900–1,000 nm (Chroma ZT640rdc-sp-UF2). Each light-pulse train duration was 1 s at 25 Hz, controlled by Arduino with 10-ms pulses. Each opto-mapping session included 45 trials (each trial is one light-pulse train), separated by around 20 s, at three light intensities (8 mW, 14 mW and 30 mW from the objective).

For optogenetic mapping sessions with inactivation of the motor thalamus, muscimol was injected using glass pipettes unilaterally into the motor thalamus of CaMKII-tTA::tetO-GCaMP6s::Vglut2-Cre triple transgenic mice: 1 mm posterior, 1 mm lateral to bregma at a depth of around 3.7 mm from the pia (approximately 0.3 μl of 5 μg μl^{-1} in cortex buffer). Mice were allowed to recover in their home cage on a heating pad for 1 h before optogenetic mapping sessions. Injected mice showed obvious turning behaviour, which confirmed successful injections.

Classification of responsive neurons. Activity ($\Delta F/F_0$) 0.5 to 1 s after the onset of a 1-s opto stimulation (LED on) was compared to activity 2 to −1 s before the onset of the opto stimulation (baseline) across all 15 trials in each light intensity. To classify a neuron as Th-excited or Th-rebound, it must be classified as such in at least one of the three light intensities according to the following criteria. For a neuron to be classified as a Th-excited neuron: first, a P value lower than 0.01 using the Wilcoxon signed-rank test when comparing LED on time to baseline. Second, an average increase in activity from baseline to LED on larger than one ($\Delta F/F_0$).

To classify a neuron as a Th-rebound neuron, the activity from the termination of the opto stimulation to 0.5 s later (LED off) was compared to the LED on activity. If a neuron had a P value lower than 0.01 using the Wilcoxon signed-rank test, when comparing LED off to LED on and an average increase in activity from LED on to LED off larger than one ($\Delta F/F_0$), this neuron was classified as a Th-rebound neuron. All neurons not classified as Th-excited or Th-rebound were classified as Th-non-responsive.

Population decoding analysis

The decoding analysis was performed using linear maximum margin classifiers (SVM) using the scikit-learn SVC package[48] to decode the time points of rewarded movements from neural data. We used the primal optimization problem formulation (dual=False) in the SVM implementation for improved computational efficiency with our high-dimensional feature space.

Preprocessing and cross-validation. The neural data (estimated spikes for soma or $\Delta F/F$ for axons) was first preprocessed. This preprocessing step consisted of binning the data to 100-ms bins. For decoding, we used tenfold cross-validation. For a given session, the trials were randomly assigned to one of the ten folds. To ensure that the neural data from nearby trials did not contaminate each other, we excluded ten time bins from the beginning and end of each trial.

Pseudopopulation generation and decoding analysis. Pseudo-simultaneous populations of neural activity ('pseudopopulations')[49] for decoding rewarded movements were created by pooling activity from neurons or axonal boutons from all mice. To discriminate between rewarded movement and non-movement time points, we first preprocessed both the neuronal and axonal activity and the behavioural data and separated them into ten cross-validation folds. Non-rewarded movement time points were excluded from the data in each fold. We excluded mice that did not have sufficient time points for both classes (rewarded and non-movement time points) in each cross-validation fold to ensure balanced sampling.

For each remaining mouse, activity from neurons or axonal boutons was *z*-scored within each set. Response vectors were assembled by randomly selecting time bins associated with rewarded movements or non-movement time points and combining these into a single *n*-length vector for *n* total units (neurons or axonal boutons) across all pooled mice. We randomly sampled 10,000 vectors and combined them into a single pseudopopulation with *n* units and 10,000 samples[50]. For each cross-validation fold, we created two such pseudopopulations: one using data from nine folds for training, and another using the held-out fold for testing.

Finally, a linear support vector classifier was trained to predict the behavioural label from the training population. For comparing decoder performance across different populations with varying numbers of units, we randomly subsampled each population from five units up to the total number of available units to ensure equal numbers before pseudopopulation generation. This subsampling was repeated 50 times to account for sampling bias. The entire pseudopopulation and decoding pipeline was repeated five times. The decoding accuracies we report are averaged across these pseudopopulation repetitions and sampling repetitions.

ShaReD

ShaReD builds on CCA, which is a statistical method used to identify and quantify relationships between two sets of variables. In the context of neuroscience, CCA can be used to uncover linear combinations of neural activity features that are maximally correlated with linear combinations of behavioural features. This allows us to investigate how neural activity patterns relate to specific aspects of behaviour.

CCA operates on two paired datasets, X (neural data) and Y (behavioural data), each containing observations across multiple time points and dimensions (e.g., neurons or behavioural features). The core objective of CCA is to find projection vectors, $\mathbf{a}$ for the neural data and $\mathbf{b}$ for the behavioural data, that maximize the correlation between the projected data points. Mathematically, this can be expressed as maximizing the following function:

$$\max \rho(\mathbf{u}, \mathbf{v}) = \max_{\mathbf{a},\mathbf{b}} \frac{\mathbf{a}^\top C_{XY} \mathbf{b}}{\sqrt{\mathbf{a}^\top C_{XX} \mathbf{a} \mathbf{b}^\top C_{YY} \mathbf{b}}},$$

where $\mathbf{u} = X\mathbf{a}$ and $\mathbf{v} = Y\mathbf{b}$ represent the projections of the neural and behavioural data onto their respective vectors, provided X and Y have been mean-centred and contain T time points. Here, $C_{XX} = \frac{X^\top X}{T}$, $C_{XY} = \frac{X^\top Y}{T}$ and $C_{YY} = \frac{Y^\top Y}{T}$ are the neural covariance, neural–behavioural cross-covariance and behavioural covariance matrices, respectively.

The optimization problem is typically solved through an eigenvalue decomposition of the covariance matrices of the data. This yields a set of canonical variates $(\mathbf{u}_1,\mathbf{v}_1)$, $(\mathbf{u}_2,\mathbf{v}_2)$, ..., each representing a pair of maximally correlated linear combinations of neural and behavioural features. The strength of the relationship between each pair is quantified by the corresponding canonical correlation coefficient (ρ).

By examining the weights within the projection vectors $\mathbf{a}$ and $\mathbf{b}$, we can gain insights into which specific neurons or behavioural features contribute most to the shared variance between the two datasets. This provides valuable information about the neural encoding of behaviour and can help us understand how neural populations represent and contribute to specific behavioural patterns.

Although CCA is effective in identifying relationships between neural and behavioural data within individuals, comparing results across multiple individuals can be challenging owing to variations in projection vectors. ShaReD addresses this limitation by finding a single set of behavioural features that are shared across all individuals and maximally correlated with the neural activity of each individual, allowing for the generalization of findings and identification of common patterns of neural–behavioural interactions.

ShaReD operates on paired neural and behavioural datasets from multiple individuals (X_1,Y_1), (X_2,Y_2), ..., (X_k,Y_k). For each individual k, the neural data matrix $X_k \in R^{T_k \times N_k}$ contains T_k time points and N_k neurons, whereas the behavioural data matrix $Y_k \in R^{T_k \times B}$ contains the same T_k time points and B behavioural features. Although the number of time points T_k may vary across individuals, the number of behavioural features B remains constant. The objective is to find a single projection vector $\mathbf{b} \in R^B$ for the behavioural data and unique neural projection vectors $\mathbf{a}_1, \mathbf{a}_2, ..., \mathbf{a}_k$ for each individual, where $\mathbf{a}_k \in R^{N_k}$. These vectors should be optimized to maximize the sum of squared correlation coefficients across all individuals, which can be achieved using a simple objective function, assuming that the neural and behavioural data have been mean-centred and that the behavioural data have been whitened (see 'Preprocessing' subsection). The complete objective function we minimize is given by

$$\begin{aligned}\min_{\mathbf{a}_k,\mathbf{b}} \mathcal{L} = &-\sum_k \omega_k (\mathbf{a}_k^\top C_{XY}^k \mathbf{b})^2 + \sum_k \alpha_k (\mathbf{a}_k^\top C_{XX}^k \mathbf{a}_k - 1) + \beta \left(\mathbf{b}^\top \frac{\sum_k C_{YY}^k}{K} \mathbf{b} - 1 \right) \\ &+ \lambda_{\text{reg}} \sum_i |\mathbf{b}_i| + \lambda_{\text{smooth}} \sum_g \sum_{j=2}^{n_g} (\mathbf{b}_j^g - \mathbf{b}_{j-1}^g)^2,\end{aligned} \tag{1}$$

where $C_{XX}^k = X_k^\top X_k / T_k$, $C_{XY}^k = X_k^\top Y_k / T_k$ and $C_{YY}^k = Y_k^\top Y_k / T_k$ are the normalized covariance matrices for each individual k with T_k time points. The average behavioural covariance matrix across individuals is defined as $\overline{C_{YY}} = \frac{1}{K} \sum_k C_{YY}^k$. In this objective function, ω_k represents a weighting factor for each individual, allowing for differential weighting of datasets on the basis of their importance or size. In our analyses, ω_k was scaled by the fraction of neurons in each dataset out of the total number of neurons across all datasets. This ensured that datasets with a larger number of neurons contributed more to the overall objective function. The Lagrange multipliers α_k and β are introduced to enforce constraints on the norms of the projection vectors, ensuring that the projected neural data for each individual ($\mathbf{a}_k^\top C_{XX}^k \mathbf{a}_k$) and the projected behavioural data ($\mathbf{b}^\top \overline{C_{YY}} \mathbf{b}$) have unit variance.

The objective function also includes regularization terms to ensure a well-behaved solution and to incorporate prior knowledge about the structure of the data. Specifically, the L1 penalty (controlled by λ_{reg}) on the behavioural projection vector $\mathbf{b}$ promotes sparsity and helps identify the most relevant and informative behavioural features that contribute most to the shared neural–behavioural correlations across individuals. Here, $\mathbf{b}_i$ corresponds to the ith element of $\mathbf{b}$. In addition, a smoothing penalty (controlled by λ_{smooth}) is applied to groups of indices that represent a discretized continuous variable, where g indexes the behavioural variable type ($g = 1$ for position, $g = 2$ for speed) and $\mathbf{b}_j^g$ denotes the projection weight for the jth lag of behavioural variable g. By enforcing smooth transitions between consecutive weights separately for position and speed features, we impose a smoothness constraint akin to temporal continuity. This incorporates prior knowledge about the temporal structure of behaviour and helps prevent overfitting to noise or idiosyncrasies in the data. For our analyses, λ_{reg} and λ_{smooth} were kept at values of 0.02 and 0.04, respectively.

We minimize the objective function in two steps. The first stage focuses on the core part of the objective function and involves solving

for the minimizer of the first three terms using an eigenvalue decomposition.

Step 1: initial optimization. To find the initial estimate for the first ShaReD component, we set $\mathbf{b}$ equal to the normalized eigenvector corresponding to the largest eigenvalue of the matrix $\sum_k \omega_k(\overline{C_{YY}})^{-1}C_{YX}^k(C_{XX}^k)^{-1}C_{XY}^k$. This eigenvector represents the initial estimate for the shared behavioural projection across all individuals. Then, for each individual k we compute the neural projection vector $\mathbf{a}_k$ as follows:

$$\mathbf{a}_k = \frac{(C_{XX}^k)^{-1}C_{XY}^k\mathbf{b}}{\sqrt{\mathbf{b}^\top C_{YX}^k(C_{XX}^k)^{-1}C_{XY}^k\mathbf{b}}}.$$

This step provides individual-specific neural projection vectors that maximize the correlation with the shared behavioural projection, serving as a starting point for further refinement. To ensure the constraints are satisfied we compute α_k and β as the following:

$$\alpha_k = \omega_k(\mathbf{a}_k^\top C_{XY}^k\mathbf{b})^2$$
$$\beta = \sum_k \omega_k(\mathbf{a}_k^\top C_{XY}^k\mathbf{b})^2$$

Step 2: Gradient descent. Owing to the inclusion of the L1 and smoothness penalties, we cannot obtain a closed-form solution for the minimizer of the entire objective function. Thus, the initial solution is refined using gradient descent to incorporate the L1 and smoothing penalties and achieve a more stable and generalizable solution. This involves iteratively updating the $\mathbf{a}_k$ and $\mathbf{b}$ vectors based on the gradients of the objective function with respect to each element. The gradient updates for each individual's neural projection vector $\mathbf{a}_k$ can be calculated as:

$$\frac{\partial\mathcal{L}}{\partial\mathbf{a}_k} = -2\omega_k(\mathbf{a}_k^\top C_{XY}^k\mathbf{b})(C_{XY}^k\mathbf{b}) + 2\alpha_k C_{XX}^k\mathbf{a}_k.$$

Similarly, the gradient update for the shared behavioural projection vector $\mathbf{b}$ is:

$$\frac{\partial\mathcal{L}}{\partial\mathbf{b}} = -2\sum_k \omega_k(\mathbf{a}_k^\top C_{XY}^k\mathbf{b})(\mathbf{a}_k^\top C_{XY}^k)^\top + 2\beta(\mathbf{b}^\top\overline{C_{YY}})^\top + \lambda_{\text{reg}}\text{sign}(\mathbf{b}) + \sum_g \nabla_{\text{smooth}}^g,$$

where ∇_{smooth}^g corresponds to the following expression, taking edge effects into account:

$$\nabla_{\text{smooth}}^g = \begin{cases} 2\lambda_{\text{smooth}}(\mathbf{b}_1^g - \mathbf{b}_2^g) & \text{for } j=1, \\ 2\lambda_{\text{smooth}}(2\mathbf{b}_j^g - \mathbf{b}_{j-1}^g - \mathbf{b}_{j+1}^g) & \text{for } 2\le j\le n_g-1, \\ 2\lambda_{\text{smooth}}(\mathbf{b}_{n_g}^g - \mathbf{b}_{n_g-1}^g) & \text{for } j=n_g. \end{cases}$$

Here, n_g represents the number of temporal lags for each behavioural variable (for example, position or speed), and $\mathbf{b}_j^g$ corresponds to the jth element of behavioural variable g. The smoothing penalty encourages temporal continuity by promoting similar weights for adjacent time lags within each behavioural variable while maintaining independence between different variables.

This iterative gradient descent process refines both the individual-specific neural projections and the shared behavioural projection to maximize the sum of squared correlations while incorporating the desired regularization properties. At each iteration of the gradient descent we recompute α_k and β to ensure the constraints are satisfied, and orthogonalize the current weights with respect to previously found components to ensure we discover independent behavioural components.

Deflation. After identifying the first ShaReD component, which consists of the shared behavioural projection vector $\mathbf{b}$ and the individual neural projection vectors $\mathbf{a}_k$, we apply deflation to remove the explained (co)variance from the data. This allows for the discovery of subsequent ShaReD components that capture additional, independent patterns of neural–behavioural correlations.

Deflation essentially removes the contribution of the identified ShaReD component from the data, forcing the subsequent optimization steps to focus on remaining sources of shared (co)variance. This is achieved by projecting the data onto a lower-dimensional subspace that is orthogonal to the space spanned by the first ShaReD component. The deflated data are then used to repeat the optimization process (steps 1 and 2) to find the next ShaReD component. This iterative process continues until a desired number of components are found or until the remaining variance of the data is negligible.

Preprocessing. Before applying ShaReD, we z-scored the data, and used singular value decomposition (SVD) to whiten the neural as well as the behavioural datasets. Given a matrix Y_k, we expressed its SVD as USV^T. To whiten the data, we removed the S matrix and calculated UV^T. This ensures that the covariance matrix is the identity matrix, which standardizes the variance across all dimensions without altering the original orientation of the data. This step ensures that all variables contribute to the analysis equally.

Projected behaviour visualization. To analyse how different movement patterns relate to the behavioural components identified by ShaReD, we first downsampled the raw lever position data from 1,000 Hz to 100 Hz by averaging within non-overlapping 10-ms windows. For each mouse, we defined the learned movement pattern by randomly selecting half of the rewarded movements and averaging their position trajectories. This sampling and averaging procedure was repeated ten times to mitigate sampling bias in identifying the learned pattern. For the remaining movements not used to define the learned pattern, we quantified their similarity to the learned movement by calculating the correlation between their lever position traces from movement onset to 1 s after onset.

We calculated lever speed from the downsampled position data by taking the absolute difference between averages of adjacent 100-ms bins (ten samples at 100 Hz). For both lever position and lever speed, we created delay matrices in which each row represented a 2.1-s sliding window (210 samples at 100 Hz), averaged into 21 non-overlapping bins of 10 samples each. These 21 bins correspond to the temporal structure of the ShaReD behavioural components (ten past lags, one present and ten future lags for each behavioural variable). The final behavioural projection was computed by applying the corresponding ShaReD weights (21 weights each for position and speed) to these binned averages, with the window sliding forward one sample (10 ms) at a time.

To visualize how different movements project onto the behavioural components, we pooled trials across all mice and grouped them on the basis of their similarity to the learned pattern. Trials were divided into high-similarity (correlation > 0.7) and low-similarity (correlation < 0.4) groups, with these thresholds chosen to ensure approximately equal numbers of trials in each group. Owing to the 2.1-s window required for the 21 ShaReD time bins (ten past lags, one present and ten future lags at 10 Hz), projecting our original movement data (−1.5 s to 2.5 s around movement onset) yielded projection values spanning from −0.45 s to 1.45 s relative to movement onset. For both high- and low-similarity groups, we calculated projections onto the behavioural components of both Th-excited and Th-non-responsive populations. The shaded regions around each trajectory represent the standard deviation across

ten iterations of randomly selecting different sets of movements to define the learned pattern.

Chemogenetic inactivation

For inactivating the motor thalamus, CNO (Enzo Life Sciences) was dissolved in sterile saline to a concentration of 2.5 mg ml^{-1} and injected i.p. at a 10 mg per kg body weight dose 45 min before the behavioural session.

For locally inactivating motor thalamic inputs in M1, CNO (1 mM) was injected into three locations in M1 (100 nl per location) at 0.5 mm below the pia surface 15 min before behavioural training.

Only mice that had hM4Di expression in the motor thalamus (VM and VAL) confirmed by histology were included in the analysis. Investigators were blinded to the identity of the mouse (hM4Di versus control) during behavioural training and to the performance of the mouse when examining hM4Di expression.

Histology

Mice were anaesthetized and transcardially perfused with ice-cold 0.1 M phosphate-buffered saline (PBS) (pH 7.4), followed by perfusion with ice-cold 4% paraformaldehyde (PFA) solution. Isolated brains were post-fixed overnight at 4 °C in 4% PFA and cryoprotected in 30% sucrose solution for at least 48 h at 4 °C. Microtome-cut (Thermo Scientific Microm HM 430) 40–50-µm free-floating brain (coronal) sections were collected in PBS and stored at 4 °C. Slices were mounted with a CC mounting medium (Sigma-Aldrich) and imaged using a fluorescence microscope (Zen and ApoTome.2, Zeiss). For visualizing axonal GCaMP6s and hM4di, coronal sections were blocked in a solution consisting of 10% donkey serum, 1% BSA and 0.3% Triton X-100 in 1× PBS for 1 h at room temperature, followed by overnight incubation at 4 °C with primary antibodies (1:1,000 chicken anti-GFP, Aves Labs (GCaMP); 1:1,000 rabbit anti-mCherry, Abcam (hM4Di)) diluted in the blocking solution. After washing, sections were incubated in secondary antibodies (1:200 goat anti-chicken Alexa Fluor 488, Jackson ImmunoResearch (GcAMP); 1:200 goat anti-rabbit, Alexa Fluor 594, Thermo Fisher Scientific (hM4Di)) for 1.5 h at room temperature. For all experiments, we verified the expression of the injection construct (axon-GCaMP6s, ChrimsonR or hM4di) in the desired brain region.

Identifying monosynaptic inputs to M1 L2/3 neurons

We analysed data from a previous study[16], representing rabies-virus-based monosynaptic retrograde labelling of direct inputs to M1 L2/3 excitatory neurons. We examined only data using Cux2-Cre and Sepw1-Cre lines, which represent putative excitatory M1 L2/3 neurons.

Statistics

For comparing decoding analyses, we used *t*-tests, with Holm–Bonferroni multiple-comparison correction where appropriate. Otherwise, non-parametric tests were used to avoid assumptions about data distributions. No statistical methods were used to predetermine sample size, but our sample sizes are similar to those reported in previous publications[1,45].

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

Supporting data are available on CodeOcean (https://codeocean.com/capsule/7424679/tree/v1). An additional dataset supporting the current work is available from the corresponding authors upon reasonable request. Source data are provided with this paper.

Code availability

The code used in this study is available on CodeOcean (https://codeocean.com/capsule/7424679/tree/v1) or from the corresponding authors upon request.

45. Peters, A. J., Lee, J., Hedrick, N. G., O'Neil, K. & Komiyama, T. Reorganization of corticospinal output during motor learning. *Nat. Neurosci.* **20**, 1133 (2017).
46. Pachitariu, M. et al. Suite2p: beyond 10,000 neurons with standard two-photon microscopy. Preprint at *bioRxiv* https://doi.org/10.1101/061507 (2017).
47. Pachitariu, M., Stringer, C. & Harris, K. D. Robustness of spike deconvolution for neuronal calcium imaging. *J. Neurosci.* **38**, 7976–7985 (2018).
48. Pedregosa, F. et al. Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
49. Boyle, L. M., Posani, L., Irfan, S., Siegelbaum, S. A. & Fusi, S. Tuned geometries of hippocampal representations meet the computational demands of social memory. *Neuron* **112**, 1358–1371 (2024).
50. Bernardi, S. et al. The geometry of abstraction in the hippocampus and prefrontal cortex. *Cell* **183**, 954–967 (2020).

Acknowledgements We dedicate this manuscript to the memory of our friend and colleague, An Wu. We thank D. Arakelyan, A. Medina, E. Hall and Y. Salazar for technical assistance; R. Gunwan for help with the FARCI algorithm; and members of the laboratory of T.K., especially N. G. Hedrick, E. Gjoni and B. P. Danskin, for discussion. This research was supported by grants from the NIH (R01 NS125298, R01 NS091010, R01 DC018545, R01 MH128746), the NSF (2024776) and the Simons Collaboration on the Global Brain Pilot Award.

Author contributions Conceptualization: A.R. and T.K. Methodology: A.R., T.K., F.H.T., M.K.B. and K.M.T. Formal analysis: A.R., F.H.T., T.K. and M.K.B. Investigation: A.R., Y.C.Y., Y.H., Q.C., X.C.W., B.C.M. and A.W. Writing (original draft): A.R. and T.K. Visualization, writing, review and editing: A.R., T.K., F.H.T. and M.K.B. Supervision and funding acquisition: T.K., M.K.B. and K.M.T.

Competing interests The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41586-025-08962-8.

Correspondence and requests for materials should be addressed to Marcus K. Benna or Takaki Komiyama.

Peer review information *Nature* thanks Christopher Rodgers and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

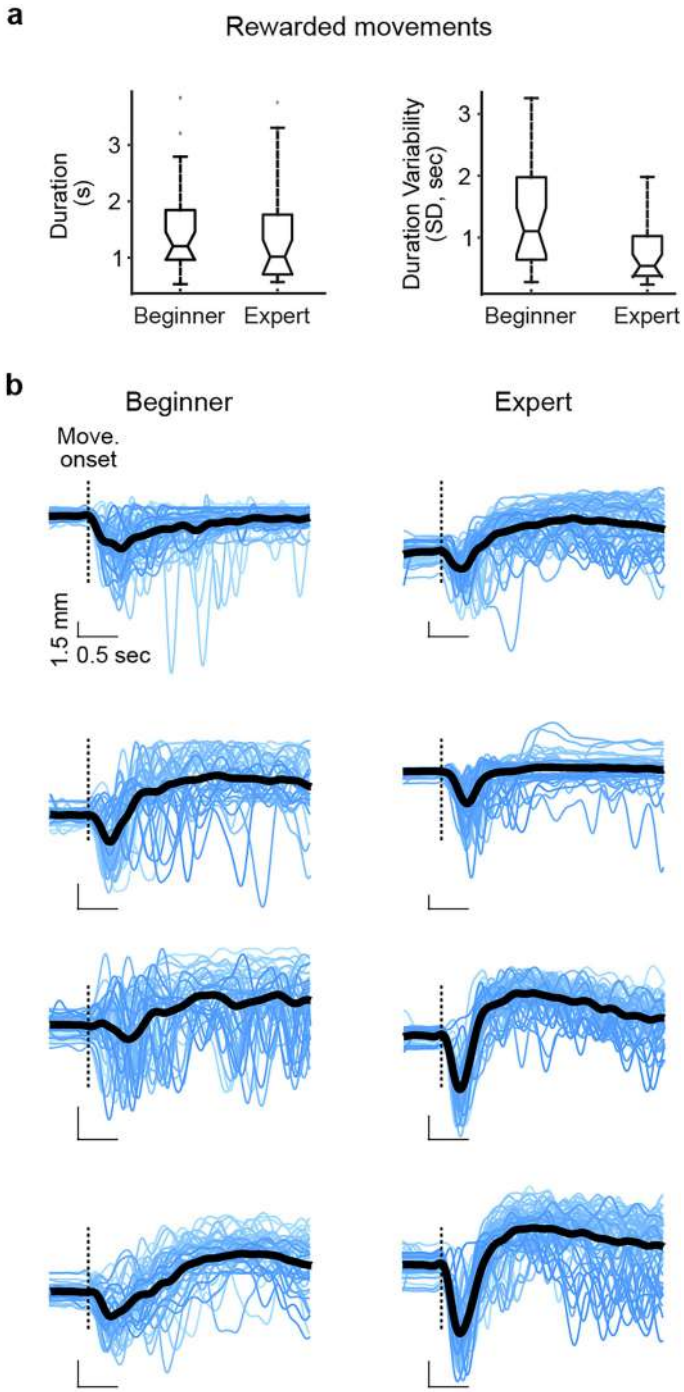

Extended Data Fig. 1 | Further analysis for rewarded movements. a, Left, the median duration of all rewarded movements across individual sessions. Right, median variability (s.d.) in rewarded movement duration across individual sessions. n = 36 for both beginner and expert sessions. **b**, Representative examples of rewarded movements. Individual (blue) and the average rewarded movements (black) during beginner (left) and expert (right) sessions.

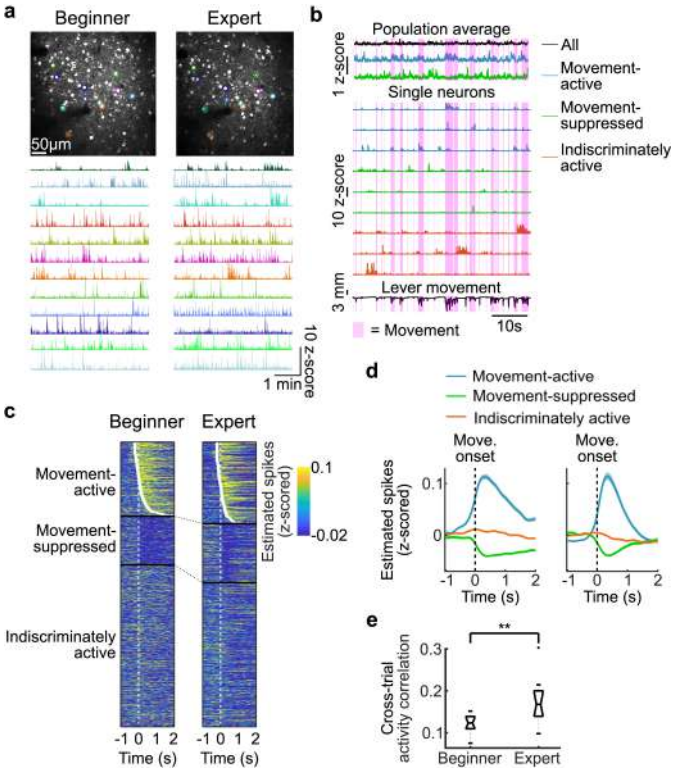

Extended Data Fig. 2 | Motor learning induces reproducible patterns of M1 L2/3 activity. a, Top, max-intensity projection of in vivo two-photon fluorescence images of GCaMP6s in M1 L2/3 of CaMKII-tTA::tetO-GCaMP6s::Vglut2-Cre triple transgenic mice, expressing GCaMP6s in cortical excitatory neurons. The same FOV imaged on the first (beginner) and last (expert) training days. Bottom, representative estimated spike traces from the same neurons. **b**, Activity from one example FOV. Top, population average of all (black), movement-active (blue), and movement-suppressed (green) neurons. Middle, representative individual neurons. Bottom, lever position. Pink highlights represent movement periods. **c**, Trial-averaged activity of movement-active (top, sorted by onset timing), movement-suppressed (middle), and indiscriminately active (bottom) neurons during beginner (left) and expert (right) sessions, aligned to the onset of rewarded movements (dashed lines). Each row represents one neuron. During beginner sessions, 30.7% of the neurons are movement-active (1132 neurons), 14.2% are movementsuppressed (525 neurons), and 55.1% are indiscriminately active (2036 neurons). During expert sessions, 35.1% of the neurons are movement-active (1035 neurons), 17.1% are movement-suppressed (504 neurons), and 47.8% are indiscriminately active (1407 neurons). **d**, Average population activity during beginner (left) and expert (right) sessions. Mean ± s.e.m. **e**, Trial-by-trial correlation of population activity patterns during rewarded movements increased over training ($p < 0.01$, rank-sum test). Box plot elements: centre line, median; box limits, upper and lower quartiles; whiskers, 1.5× interquartile range; points, outliers. For **b**,**c**,**d**,**g**–**i**, n = 7 mice; 14 beginner sessions, 13 expert sessions. All rank-sum tests are two-sided.

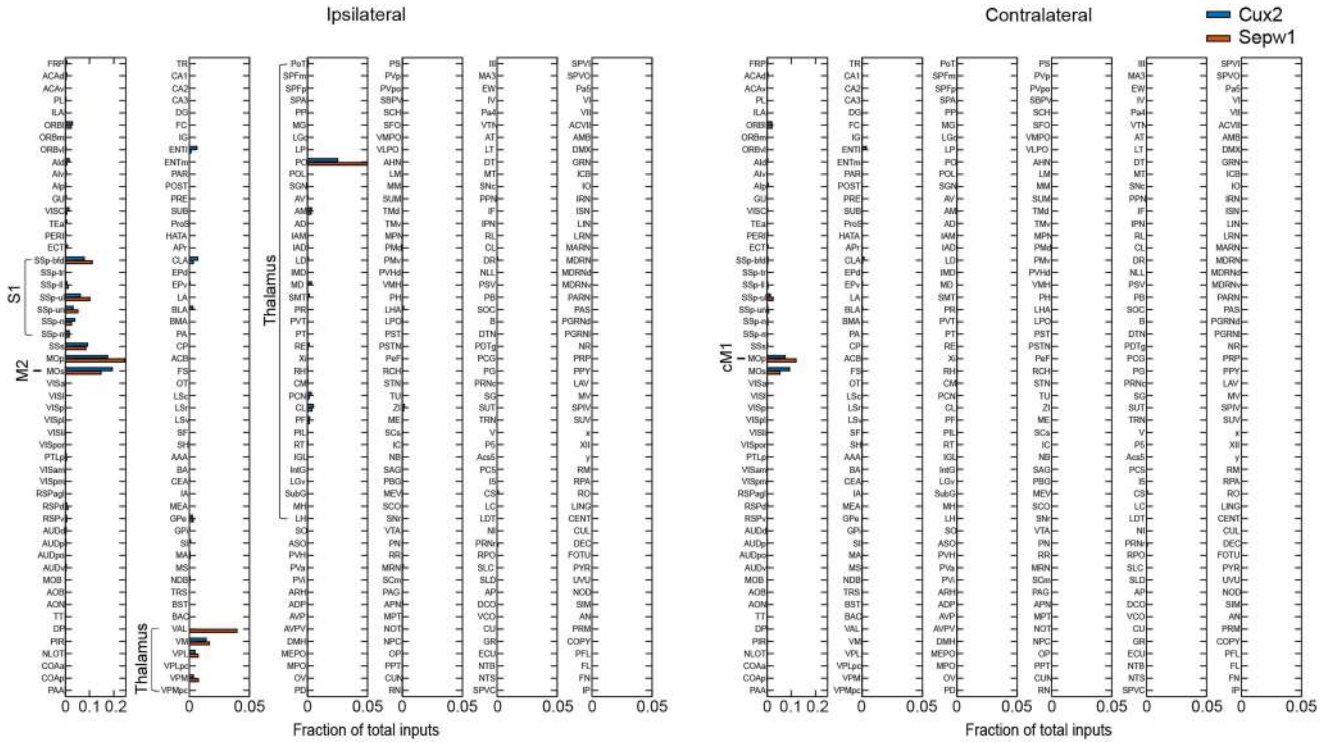

Extended Data Fig. 3 | Monosynaptic inputs to M1 L2/3 neurons. Summary of rabies-virus-based monosynaptic retrograde labelling of direct inputs to M1 L2/3 excitatory neurons. Cux2 and Sepw1 Cre lines were used to limit the starter cells to L2/3 excitatory neurons. n = 2 per Cre line. Data from a previous study[16]. MOp, primary motor cortex; SSp, primary somatosensory cortex; SSs, secondary somatosensory cortex; VAL, ventral anterior-lateral complex of the thalamus; VM, ventral medial nucleus of the thalamus; VPL, ventral posterolateral nucleus of the thalamus; VPM, ventral posteromedial nucleus of the thalamus; PO, posterior complex of the thalamus.

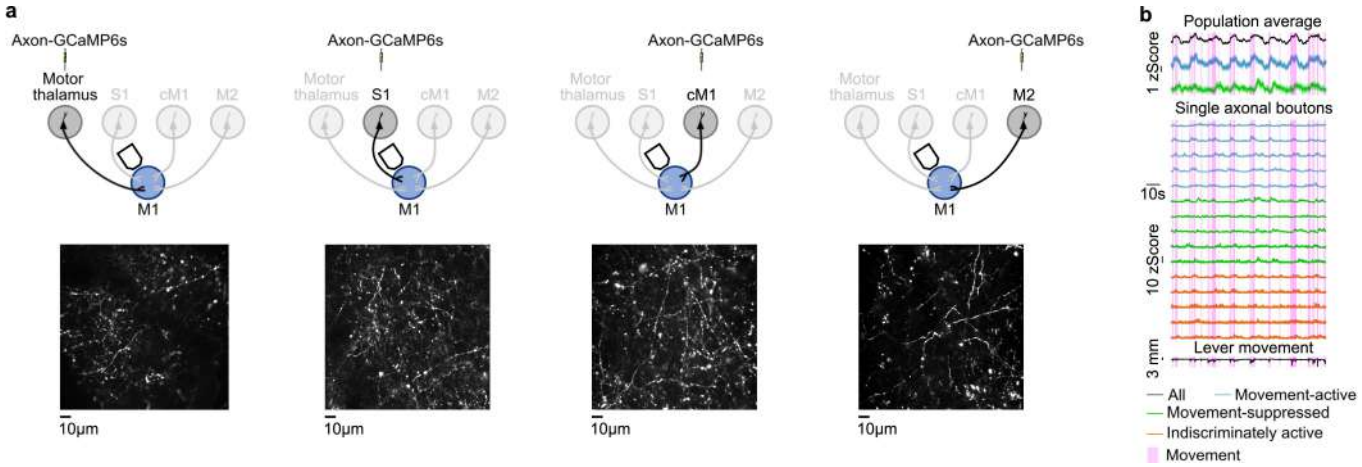

Extended Data Fig. 4 | Projection-specific functional imaging. Related to Fig. 1. **a**, Top, schematic of injections to express axon-GCaMP6s in different input areas and imaging their axons innervating M1. Bottom, max-intensity projection of in vivo two-photon fluorescence images of axon-GCaMP6s expressed in the axons in M1 from each input source. **b**, Top, example activity from corresponding FOVs in a., showing the population average of all (black), movement-active (blue), and movement-suppressed (green) axonal boutons. Middle: representative individual axonal boutons. Bottom, lever position. Pink highlights represent movement periods.

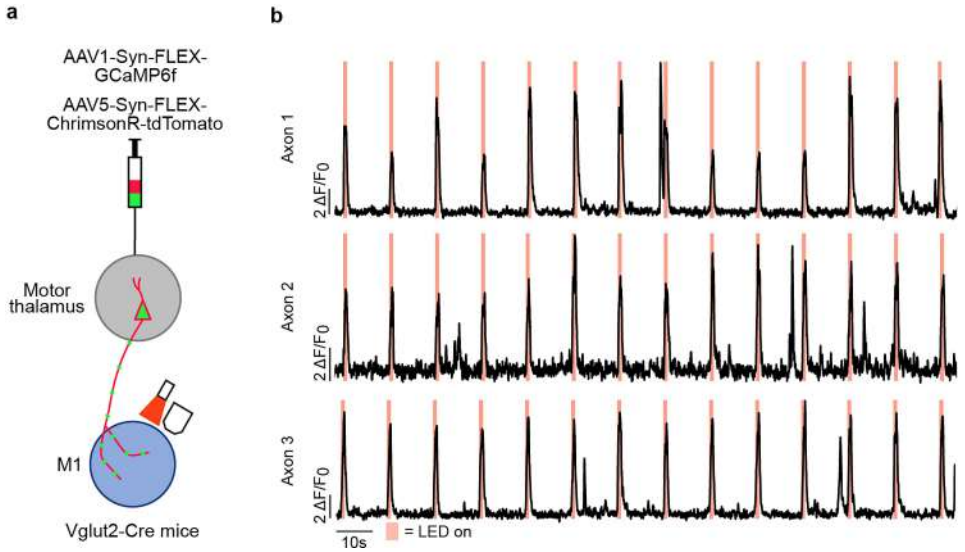

Extended Data Fig. 5 | Optogenetic stimulation of thalamic axons in M1. **a**, Schematic for experiments to validate optogenetic stimulation of thalamic inputs innervating M1 by expressing ChrimsonR and GCaMP6f in thalamic neurons and exciting and imaging their axons in M1. **b**, Three example axonal segments during optogenetic stimulation showing a reliable excitation after the stimulation of thalamic inputs (orange vertical bars).

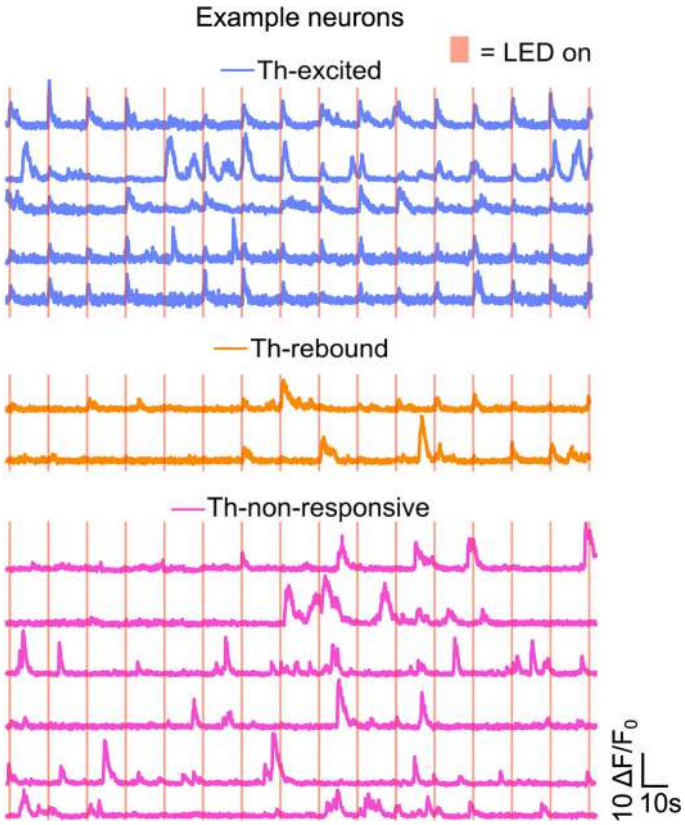

Extended Data Fig. 6 | Representative neurons during optogenetic stimulation. Responses of example M1 L2/3 neurons from a single FOV to the optogenetic stimulation of thalamic inputs. Orange vertical bars represent stimulation periods.

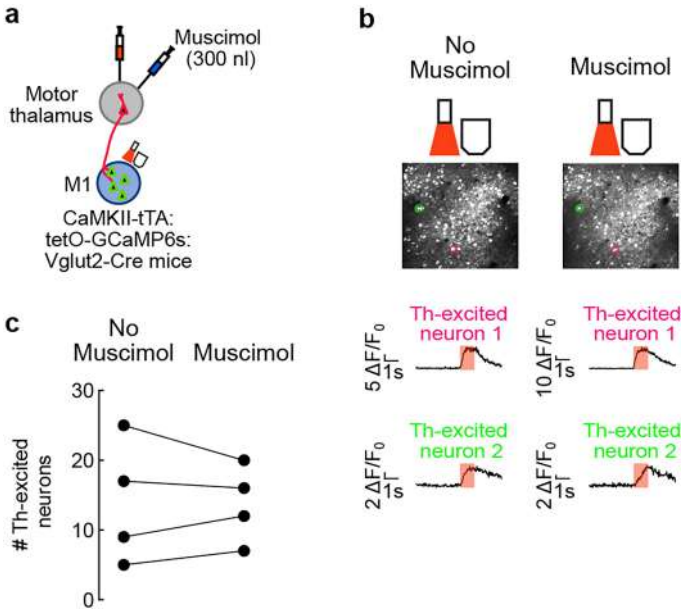

Extended Data Fig. 7 | Optogenetic mapping with motor-thalamus inactivation. a, Experimental strategy. ChrimsonR is expressed in thalamic neurons, and imaging and optogenetic stimulation are done in M1 before and after muscimol injection into the motor thalamus. **b**, Example response to thalamic axon stimulation from two Th-excited neurons before (left) and after (right) muscimol injection into the motor thalamus. **c**, Total number of Th-excited neurons per FOV before (left) and after (right) muscimol injection into the motor thalamus. Each line is one FOV (n = 4 FOVs from 2 mice).

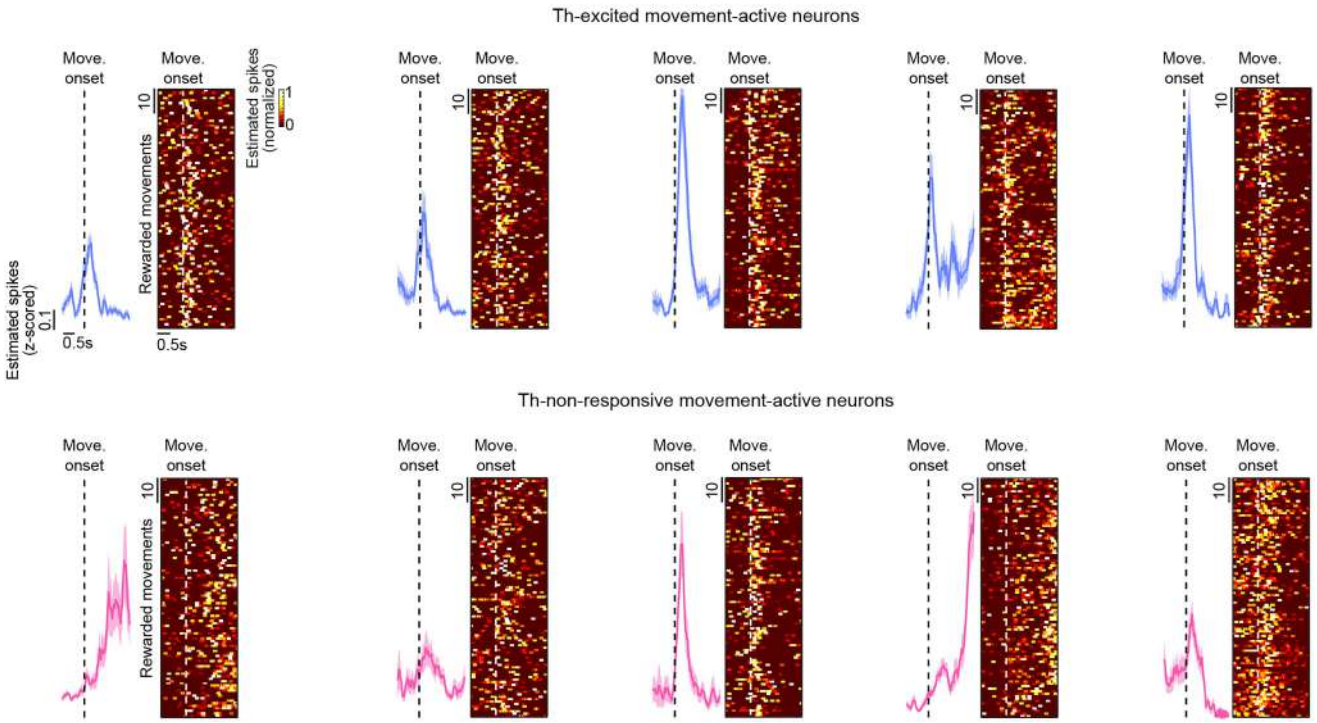

Extended Data Fig. 8 | Examples of movement-active neurons. The activity of example Th-excited (top) and Th-non-responsive (bottom) movement-active neurons aligned to the onset of rewarded movements (dashed lines). Left, mean ± s.e.m. Right, heat maps of individual trials.

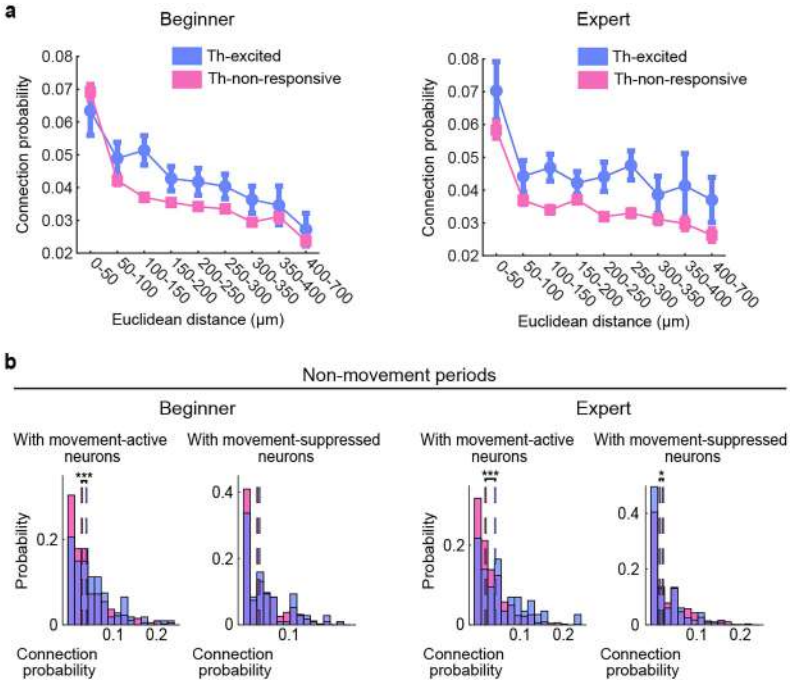

Extended Data Fig. 9 | Functional connectivity probability. a, Functional connection probability as a function of distance between neurons for Th-excited (blue) and Th-non-responsive (pink) neurons during beginner (left) and expert (right) sessions. **b**, Correlation-based functional connectivity measured while excluding all movement periods. The results are similar to the functional connectivity measured using all time periods. Left, in beginners, Th-excited neurons have higher connection probability with movement-active neurons than Th-non-responsive neurons with movement-active neurons ($p < 0.001$, signed-rank test). Connection probabilities with movement-suppressed neurons are similar between Th-excited and Th-non-responsive neurons ($p = 0.1260$). Right, in experts, Th-excited neurons have a higher connection probability with movement-active neurons and a lower connection probability with movement-suppressed neurons than Th-non-responsive neurons ($p < 0.001$ and 0.05, signed-rank test). All tests are two-sided.

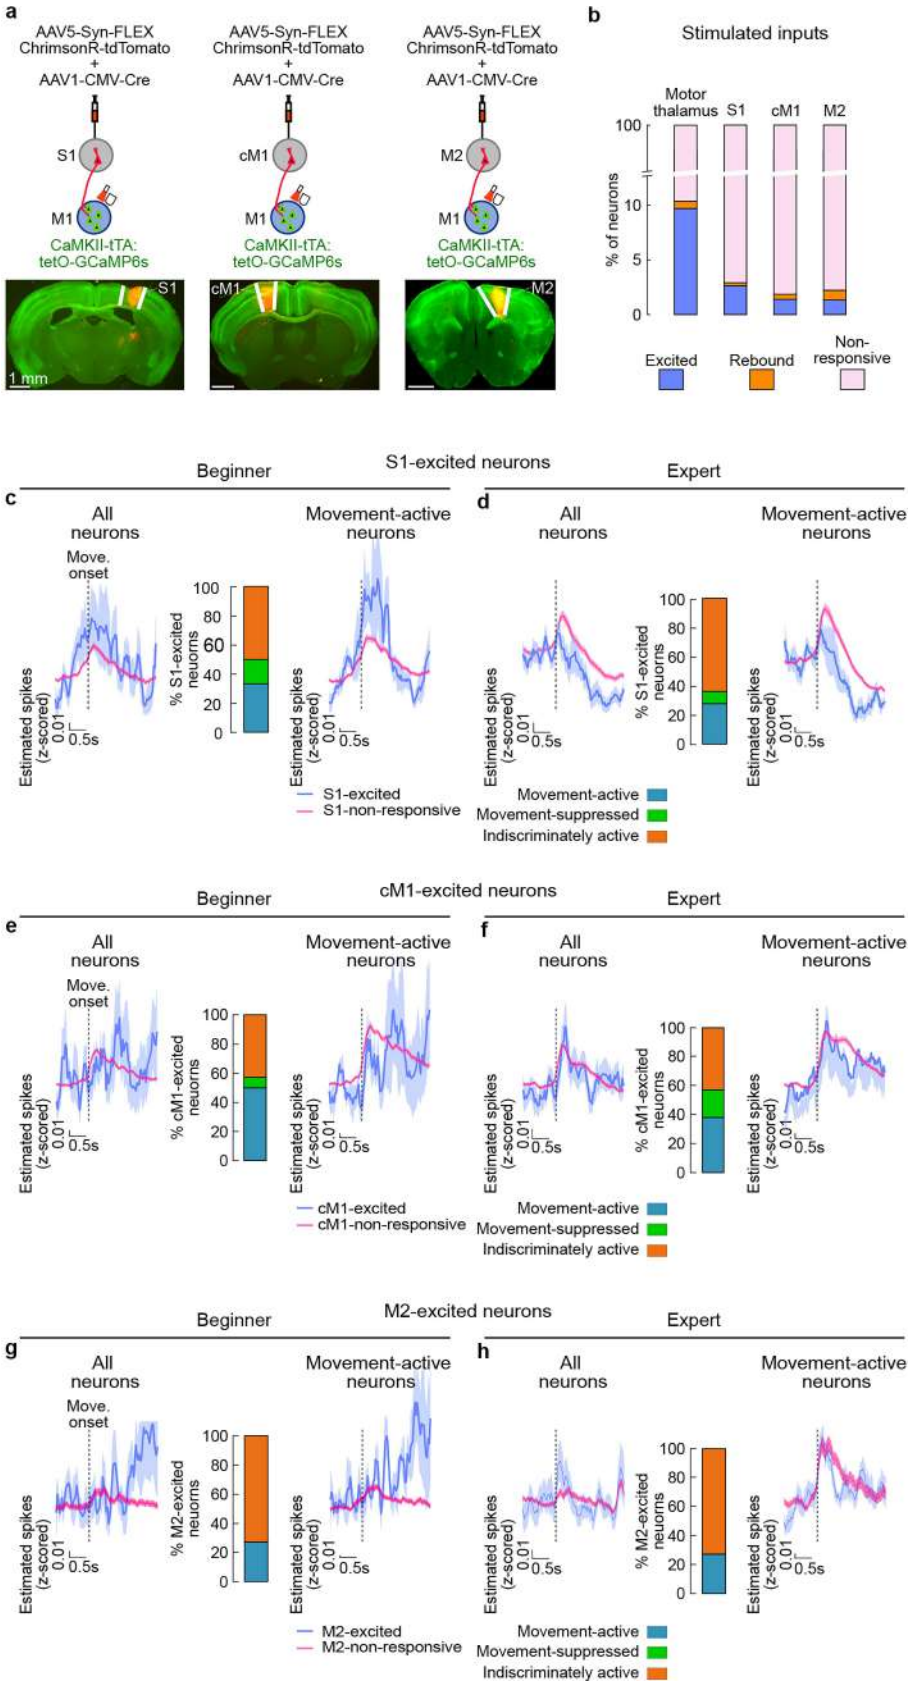

Extended Data Fig. 10 | See next page for caption.

Extended Data Fig. 10 | S1-excited, cM1-excited and M2-excited M1 L2/3 neurons do not preferentially encode the learned movement. a, Top, experimental strategy. ChrimsonR is expressed in S1, cM1, and M2 neurons, and imaging and optogenetic stimulation are performed in M1. Bottom, coronal sections showing the expression of ChrimsonR-tdTomato neurons (red) and GCaMP6s in cortical neurons (green). **b**, Fractions of excited, rebound, and non-responsive neurons by stimulation of the 4 inputs during the opto-mapping session 0 (motor thalamus, n = 1494 neurons from 7 mice; S1, n = 754 neurons from 5 mice; cM1, n = 1288 neurons from 5 mice; M2, n = 1792 neurons from 4 mice). **c**, Left, the trial-average activity of all S1-excited and S1-non-responsive (n = 856 neurons from 5 mice) neurons from beginner sessions aligned to rewarded movement onset (dashed line). Mean ± s.e.m. Middle, fractions of movement-active, movement-suppressed, and indiscriminately active neurons within the S1-excited population in beginners. Right, same as the left panel but for movement-active neurons. **d**, Same as **c** but for experts (n = 760 neurons from 5 mice). **e**, Same as **c** but for cM1-excited neurons (n = 1572 neurons from 5 mice). **f**, Same as **d** but for cM1-excited neurons (n = 1203 neurons from 5 mice). **g**, Same as **c** but for M2-excited neurons (n = 331 neurons from one mouse). **h**, Same as **d** but for M2-excited neurons (n = 493 neurons from one mouse).

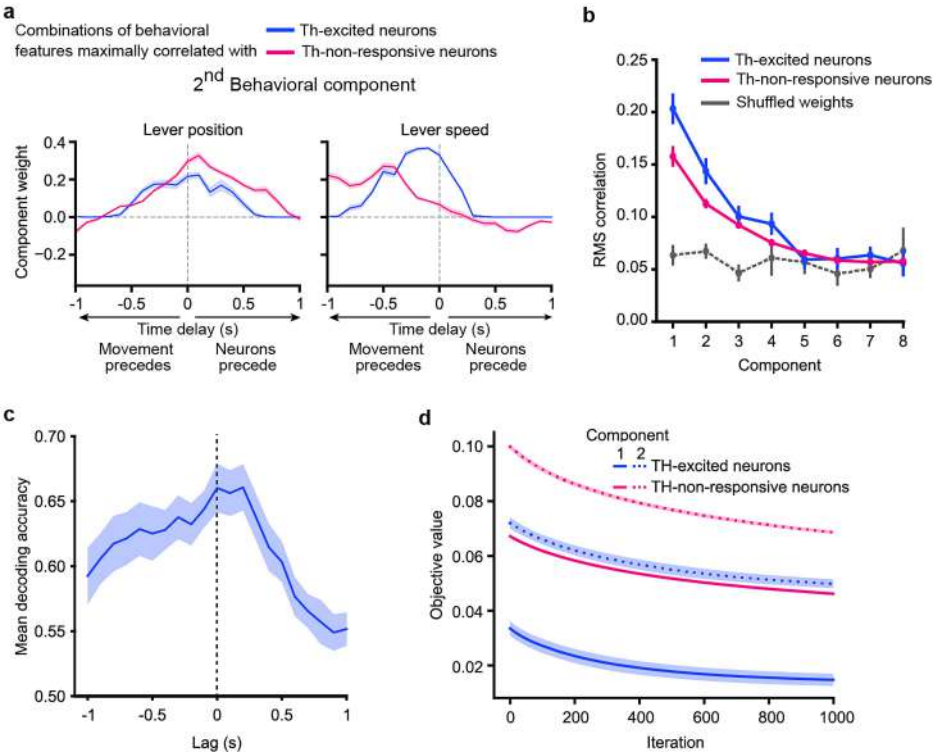

Extended Data Fig. 11 | Further analysis related to ShaReD. a, Weights (with unit L2-norm) for the second ShaReD behavioural component for movement-active Th-excited neurons (blue) or movement-active Th-non-responsive neurons (pink). Shaded areas show 95% confidence intervals of the mean over cross-validation folds. For Th-non-responsive neurons, the weights are computed 50 times for different samplings of neurons to match the number of Th-excited neurons and averaged over samples. **b**, Cross-validated measure of the root mean square correlation (between projected behaviour and projected neural activity) for the first 8 ShaReD components for movement-active Th-excited neurons (blue) and movement-active Th-non-responsive neurons (pink). Mean with 95% confidence intervals over cross-validation folds. **c**, Cross-validated decoding accuracy for the first behavioural component for movement-active Th-excited neurons when the behavioural projection weights are time-shifted. The decoding accuracy is the highest around zero time-shift (lag), i.e., when neural activity precedes movements as in Fig. 3b. Shaded areas show 95% confidence intervals of the mean over cross-validation folds. **d**, Convergence of the ShaReD objective function across 1,000 gradient descent iterations, shown for both first (solid lines) and second (dotted lines) components of Th-excited (blue) and Th-non-responsive (pink) neurons. Shaded areas show 95% confidence intervals of the mean over cross-validation folds.

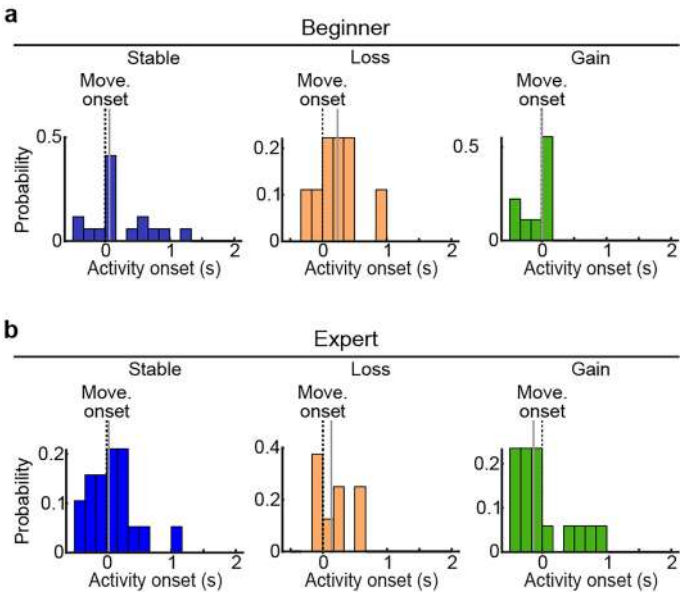

Extended Data Fig. 12 | Activity onset of different Th-excited populations. **a**,**b**, Histograms of the activity onset of individual movement-active neurons relative to the onset of rewarded movements during beginner (**a**) and expert (**b**) sessions. Vertical grey lines indicate medians.

nature portfolio

Corresponding author(s): Takaki Komiyama, Marcus Benna

Last updated by author(s): 03/17/2025

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our Editorial Policies and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed
☐	☒ The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒ A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒ The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☐	☒ A description of all covariates tested
☐	☒ A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒ A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒ For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted *Give P values as exact values whenever suitable.*
☒	☐ For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☒	☐ For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☐	☒ Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	Two-photon imaging was performed using ScanImage (v5) software running on MATLAB (2018a, Mathworks). Behavioral data was acquired via the continuous monitoring of lever position via a LabJack acquisition device and Ephus software running on MATLAB (version 2011b, Mathworks) working with custom code running on LabVIEW (version 9.0, National Instruments) to monitor threshold crossing. The behavioral setup was controlled by MATLAB (version 2010, Mathworks). Fluorescent images for histology were acquired using Zen2 (blue edition) software.
Data analysis	Analysis of two-photon data was performed using a combination of previously published code from the lab and additional custom code to fit the needs of the current experiments (all in MATLAB). Pseudopopulation generation, decoding analysis and Shared Representation Discovery (ShaReD) was done using Python.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our policy

The dataset supporting the current work is available from the corresponding authors upon reasonable request. Source data are provided with this paper. Sourced data for Figure 5 and EDF8 is on CodeOcean

Research involving human participants, their data, or biological material

Policy information about studies with human participants or human data. See also policy information about sex, gender (identity/presentation), and sexual orientation and race, ethnicity and racism.

Reporting on sex and gender	*Use the terms sex (biological attribute) and gender (shaped by social and cultural circumstances) carefully in order to avoid confusing both terms. Indicate if findings apply to only one sex or gender; describe whether sex and gender were considered in study design; whether sex and/or gender was determined based on self-reporting or assigned and methods used. Provide in the source data disaggregated sex and gender data, where this information has been collected, and if consent has been obtained for sharing of individual-level data; provide overall numbers in this Reporting Summary. Please state if this information has not been collected. Report sex- and gender-based analyses where performed, justify reasons for lack of sex- and gender-based analysis.*
Reporting on race, ethnicity, or other socially relevant groupings	*Please specify the socially constructed or socially relevant categorization variable(s) used in your manuscript and explain why they were used. Please note that such variables should not be used as proxies for other socially constructed/relevant variables (for example, race or ethnicity should not be used as a proxy for socioeconomic status). Provide clear definitions of the relevant terms used, how they were provided (by the participants/respondents, the researchers, or third parties), and the method(s) used to classify people into the different categories (e.g. self-report, census or administrative data, social media data, etc.) Please provide details about how you controlled for confounding variables in your analyses.*
Population characteristics	*Describe the covariate-relevant population characteristics of the human research participants (e.g. age, genotypic information, past and current diagnosis and treatment categories). If you filled out the behavioural & social sciences study design questions and have nothing to add here, write "See above."*
Recruitment	*Describe how participants were recruited. Outline any potential self-selection bias or other biases that may be present and how these are likely to impact results.*
Ethics oversight	*Identify the organization(s) that approved the study protocol.*

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☒ Life sciences ☐ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	No statistical methods were used to pre-determine sample sizes; however, our sample sizes are comparable to those reported in previous studies
Data exclusions	We excluded mice that did not significantly improve their performance over learning. For the chemogenetic inactivation experiment, animals with minimal hM4Di expression within the thalamus or high expression outside the thalamus were excluded (see Blinding).
Replication	Experiments and simulations were replicated as follows: - For the experiments described in Figures 3,4,5 and 6 - 10 distinct recordings per learning stage were used (7 mice, 1 or 2 fields of view per mouse). - For the experiment described in Figure 2 - Multiple distinct recordings per learning stage were used (taken from 7, 4, 7 and 7 mice for thalamic, S1, cM1 and M2 inputs, respectively). - For the experiment described in Figures 7 - 9 (hM4Di group) and 6 (tdTomato group) mice were used.

Randomization	Animals used in this study were not selected based on prerequisite features other than general animal well-being for allocation into a
Blinding	For the chemogenetic inactivation experiment, Investigators were blinded to the animal's performance when training and examining hM4Di expression.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☐	☒ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology and archaeology
☐	☒ Animals and other organisms
☒	☐ Clinical data
☒	☐ Dual use research of concern
☒	☐ Plants

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

Antibodies

Antibodies used	Chicken anti-GFP, Aves Labs; Goat anti-chicken 488, Jackson ImmunoResearch; Rabbit anti-mCherry, Abcam; Goat anti-Rabbit, Alexa Fluor 594, Thermo Fisher Scientific.
Validation	Chicken anti-GFP, Aves Labs validated in immunofluorescence here - https://www.aveslabs.com/products/anti-green-fluorescent-protein-antibody-gfp#IHC. Goat anti-chicken 488, Jackson ImmunoResearch validated in immunofluorescence here - https://www.citeab.com/antibodies/2034477-103-545-155-alexa-fluor-488-affinipure-goat-anti-c?utm_campaign=Widget+All+Citations&utm_medium=Widget&utm_source=Jackson+Immunoresearch&utm_term=Jackson+ImmunoResearch; Rabbit anti-mCherry, Abcam validated in immunofluorescence here - https://www.abcam.com/en-us/products/primary-antibodies/mcherry-antibody-ab167453?srsltid=AfmBOoo9F0T3qXB9B6-E9_tVDHGe76oR3U5wLcO1T8BsD1gekY9QWOR9#; Goat anti-Rabbit, Alexa Fluor 594, Thermo Fisher Scientific validated in immunofluorescence here - https://www.thermofisher.com/antibody/product/Goat-anti-Rabbit-IgG-H-L-Cross-Adsorbed-Secondary-Antibody-Polyclonal/A-11012

Animals and other research organisms

Policy information about studies involving animals; ARRIVE guidelines recommended for reporting animal research, and Sex and Gender in Research

Laboratory animals	Mice - C57BL/6 (Carles River), Vglut2-Cre (Vglut2-ires-cre, JAX stock #016963), CaMKII-tTA (B6;CBA-Tg(Camk2a-tTA)1Mmay/J, JAX stock #003010), and tetO-GCaMP6s (B6;DBA-Tg(tetO-GCaMP6s)2Niell/J, JAX stock #024742)
Wild animals	*Provide details on animals observed in or captured in the field; report species and age where possible. Describe how animals were caught and transported and what happened to captive animals after the study (if killed, explain why and describe method; if released, say where and when) OR state that the study did not involve wild animals.*
Reporting on sex	Both males and females were used in the study.
Field-collected samples	*For laboratory work with field-collected samples, describe all relevant parameters such as housing, maintenance, temperature, photoperiod and end-of-experiment protocol OR state that the study did not involve samples collected from the field.*
Ethics oversight	All animal procedures were performed in accordance with guidelines and protocols approved by the UCSD Institutional Animal Care and Use Committee and the National Institutes of Health (NIH).

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Plants

Seed stocks

Report on the source of all seed stocks or other plant material used. If applicable, state the seed stock centre and catalogue number. If plant specimens were collected from the field, describe the collection location, date and sampling procedures.

Novel plant genotypes

Describe the methods by which all novel plant genotypes were produced. This includes those generated by transgenic approaches, gene editing, chemical/radiation-based mutagenesis and hybridization. For transgenic lines, describe the transformation method, the number of independent lines analyzed and the generation upon which experiments were performed. For gene-edited lines, describe the editor used, the endogenous sequence targeted for editing, the targeting guide RNA sequence (if applicable) and how the editor was applied.

Authentication

Describe any authentication procedures for each seed stock used or novel genotype generated. Describe any experiments used to assess the effect of a mutation and, where applicable, how potential secondary effects (e.g. second site T-DNA insertions, mosiacism, off-target gene editing) were examined.