NEUROSCIENCE

Disengagement of motor cortex from movement control during long-term learning

Eun Jung Hwang*, Jeffrey E. Dahlen*, Yvonne Yuling Hu, Karina Aguilar, Bin Yu, Madan Mukundan, Akinori Mitani, Takaki Komiyama†

Motor learning involves reorganization of the primary motor cortex (M1). However, it remains unclear how the involvement of M1 in movement control changes during long-term learning. To address this, we trained mice in a forelimb-based motor task over months and performed optogenetic inactivation and two-photon calcium imaging in M1 during the long-term training. We found that M1 inactivation impaired the forelimb movements in the early and middle stages, but not in the late stage, indicating that the movements that initially required M1 became independent of M1. As previously shown, M1 population activity became more consistent across trials from the early to middle stage while task performance rapidly improved. However, from the middle to late stage, M1 population activity became again variable despite consistent expert behaviors. This later decline in activity consistency suggests dissociation between M1 and movements. These findings suggest that long-term motor learning can disengage M1 from movement control.

INTRODUCTION

Motor learning is supported by reorganization of motor circuits. Central to this process is the primary motor cortex (M1), where many types of changes have been described during motor learning (1, 2). For example, longitudinal imaging studies have established the high degree of spatiotemporal specificity of the formation and elimination of dendritic spines in M1 during motor learning (3–7). In addition to structural changes, it has been shown that synaptic plasticity such as long-term potentiation occurs in M1 during motor learning (8). These synaptic changes likely contribute to changes in the activity pattern of M1, leading to an improvement in the generated movements. Longitudinal recordings of M1 neural ensembles have consistently revealed that motor learning induces changes in the spatiotemporal activity pattern of M1 neurons during the production of learned movements (6, 9–12).

These previous studies largely focused on the early stage of learning during which behavioral performance rapidly improves. However, motor learning evolves through multiple stages (13, 14). The early stage that accompanies a rapid and overt improvement of performance is followed by a later stage during which extensive training may not result in obvious improvements in generated movements but can lead to a more effortless execution of the learned movement (13). It remains unknown how the later stage of learning affects the involvement of M1 in movement control. Here, we consider two possible scenarios during these two stages of motor learning. For one, it is possible that the M1 activity pattern that is acquired during the early stage as described above is maintained, and the stable performance in the later stage remains under M1 control and involves a stable M1 activity pattern. Alternatively, the later stage may involve additional changes in M1 activity and/or changes in M1 contribution to movement control, although there are little overt changes in behavior.

Here, we examined the involvement of M1 in movement control during motor learning and prolonged training over months by ap-plying optogenetic inactivation and two-photon calcium imaging during long-term training in a forelimb motor task in mice. We found that the exact same inactivation protocol had markedly different effects on movements depending on the learning stage. Furthermore, longitudinal imaging uncovered a biphasic evolution of M1 activity patterns. The change during the later phase occurred, although the expert-level performance was maintained. These results indicate that the involvement of M1 in movements dynamically changes over the course of long-term learning.

RESULTS

To investigate the involvement of M1 in movement control throughout initial motor learning and subsequent prolonged training, we trained mice in a joystick press task daily for 60 days (n = 12 mice). In this task, head-fixed mice grabbed a joystick with their left forepaw and were required to press it into a target zone in a two-dimensional space after the auditory go cue to receive a water reward (100 trials per daily session; Fig. 1A and Materials and Methods). As training progressed, the success rate (i.e., the fraction of successful trials) increased rapidly within the first few sessions and stabilized at 95 ± 5.6% (mean ± SD) after session 20 (Fig. 1B). Furthermore, the time required to acquire a reward in successful trials (i.e., duration from the go cue to target entry) decreased over training (Fig. 1B). Therefore, over time, mice learned to attain the behavioral goal more reliably and rapidly in the joystick task, important criteria for motor learning.

In addition to the improvement in the task performance as described above, daily training led to an increased stereotypy of movements, which is another sign of motor learning (15). For example, the movement onset time became more consistent across trials, and the duration from movement onset to target entry also became more regular (Fig. 1C). Furthermore, the trajectories of movements, measured by the two-dimensional joystick position, became more similar from trial to trial within and across days (Fig. 1, D to F). Thus, both the task performance and the stereotypy of task-relevant kinematics improved with training in the joystick task, indicating motor learning.
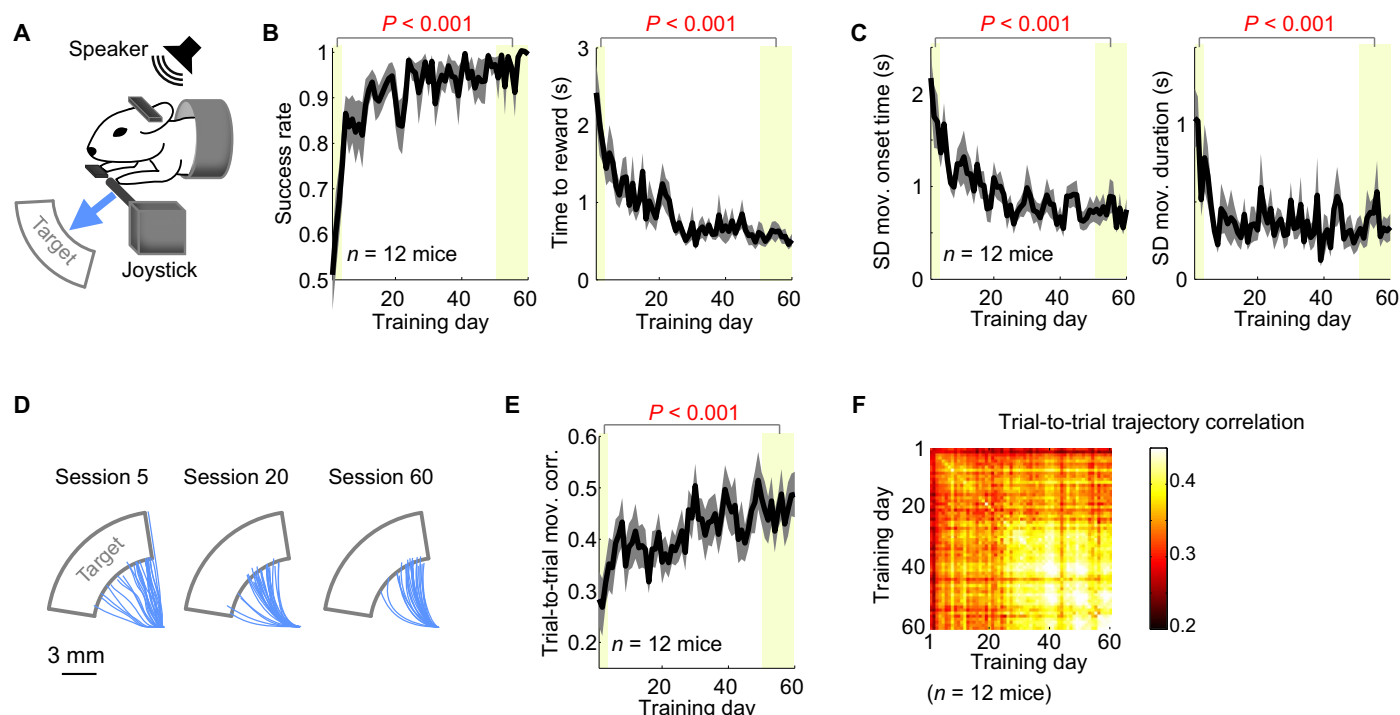
The task performance and movement consistency measures did not improve at a constant rate over the 60-day training. Instead, they rapidly improved in the first phase of training, and the improvement

Neurobiology Section, Center for Neural Circuits and Behavior, Department of Neurosciences, and Halıcıoğlu Data Science Institute, University of California, San Diego, La Jolla, CA, USA.
*These authors contributed equally to this work.
†Corresponding author. Email: tkomiyama@ucsd.edu

Fig. 1. Task performance and movement consistency improve over long-term training in the joystick task. (**A**) The joystick task setup. The mouse is required to move the joystick into the target upon the auditory go cue to receive a water reward. (**B**) The success rate (i.e., fraction of trials that acquired reward over all trials; left) and time required to acquire reward (right) as a function of training day. Mean ± SEM ($n = 12$ mice). P values are from two-sided pairwise comparison between the average values in the early (days 1 to 3) and late learning stages (days 50 to 60). (**C**) Variability of movement onset time (left) and movement duration (right) as a function of training day. SDs were used to measure variability. (**D**) Joystick movement trajectories from three different training days of a single mouse. Thirty trials in each condition are shown. (**E**) Trial-to-trial movement trajectory correlation as a function of training day. (**F**) Trial-to-trial trajectory correlation between two training days, averaged across 12 mice. The diagonal squares represent the trial-to-trial correlation within single training days plotted in (E).

was maintained in the subsequent phase. We defined the expert level of behavior, separately for each mouse, based on their mean (μ) and SD (σ) of each measure of performance and consistency in the final 10 training days. When all five measures (i.e., success rate, time to reward, movement onset time variability, movement duration variability, and trial-to-trial movement correlation) reached within 1σ from μ for three consecutive days, we declared that the mouse reached its expert level. We found that the expert level was achieved in 23 ± 4 days (mean ± SEM) across 12 mice.

It has been shown that, in some conditions, prolonged training can transform goal-directed movements into habitual responses, such that the responses after long-term training become less sensitive to the availability or value of reward (16). To test whether the joystick press movements became habitual after the 60-day training, we performed two additional experiments. First, after 60 days of training, mice were placed on ad lib water for 5 days to devalue the water reward. These satiated mice did not perform the task (fig. S1A). Second, during the task performance of the expert mice, we omitted the reward in 20 successive trials. During these reward omission trials, mice quickly reduced the rate of responses (fig. S1B). These results indicate that, despite the long-term training and high level of performance, the joystick press movements in our task remained goal directed even after 60 days of training.

M1-dependent movements can become M1 independent after prolonged training

To directly examine the involvement of M1 in movement control at different learning stages, we compared the effects of M1 inactivation

on joystick press movements at three different learning stages: early, middle, and late. Early inactivation was performed during days 4 to 8, middle inactivation during days 20 to 25, and late inactivation during days 61 to 69 after 60 days of training (Fig. 2A). The middle stage was chosen to overlap with the average time period in which mice reached their expert level as described above (day 23 ± 4). Inactivation was induced by optogenetically activating parvalbumin (PV)–positive inhibitory neurons (Materials and Methods; fig. S2) (6, 17). PV neurons in M1 were activated by blue light-emitting diode (LED) light directed into the bilateral cranial windows over M1 on inactivation days (Fig. 2B). Using this technique, we inactivated M1 from the trial onset in a subset (~12%) of randomly interleaved trials. Although the cortical control of forelimb movements is believed to be driven predominantly by the contralateral hemisphere, it has been suggested that the ipsilateral side is also involved (18), especially when the contralateral side is lesioned (19, 20). Thus, to examine the necessity of M1 as a whole, we applied bilateral inactivation. To address any nonspecific light effects, we interleaved "head-bar control days" in which the lights were directed at the head bar away from the cranial windows (Fig. 2B). Behaviors in the light-off trials were generally equivalent between the head-bar control and M1 inactivation days in all learning stages (fig. S3). In the analyses hereafter, we compared behaviors in light-on trials between head-bar control and M1 inactivation days to assess M1 inactivation effects.

Inactivation in the early stage of training ($n = 13$ mice) resulted in a severe impairment in the performance of the joystick task as indicated by the significantly reduced success rate (Fig. 2, C and D). We further characterized the impairment by categorizing failed trials into two types:
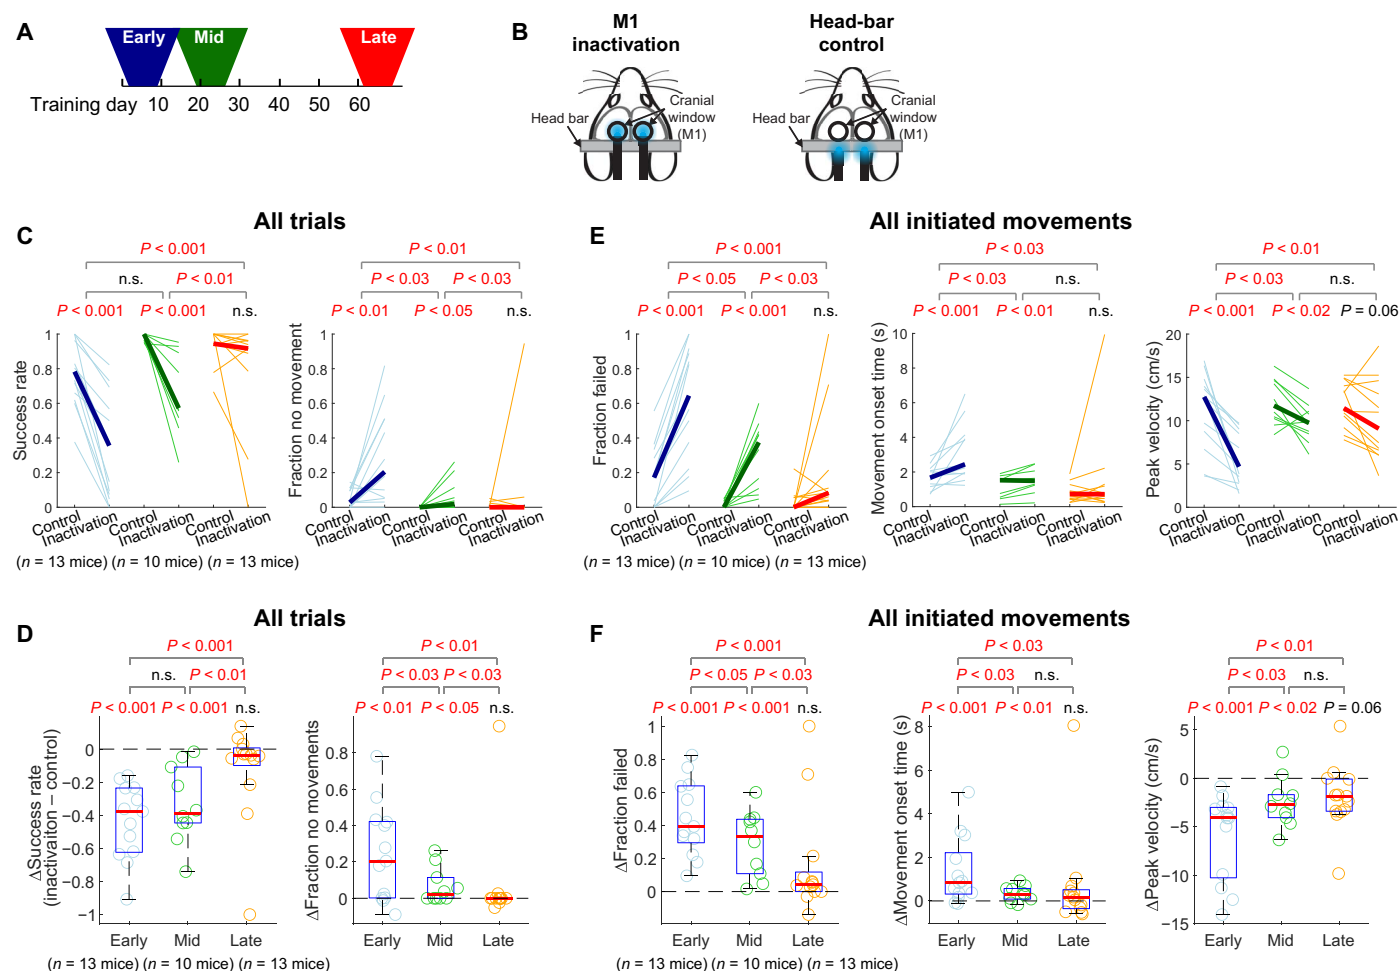
Fig. 2. M1 inactivation effects on movements gradually decline during long-term training. (**A**) Inactivation/head-bar control experiments in the early (days 1 to 9), mid (days 19 to 26), or late learning stage (days 61 to 69). (**B**) M1 inactivation and head-bar control days were randomly interleaved in each learning stage. The blue LED light was turned on in randomly selected trials (~12%) in each day. These light-on trials in the inactivation and head-bar days are referred to as inactivation and control trials, respectively, in the following comparisons. (**C**) The success rate and the fraction of trials in which mice made no movements out of all trials, in control versus inactivation trials, at the early, middle, and late learning stage. Thin lines represent individual mice, and thick lines represent medians across mice. Two-sided pairwise comparison between control and inactivation trials within each learning stage are displayed. For the effect size comparison between different learning stages, one-sided unpaired comparison was performed on the differences between control and inactivation trials in each stage. n.s., not significant. (**D**) Inactivation-induced changes (inactivation − control) in the success rate and the fraction of no movements out of all trials. The circles represent individual mice. The edges of the boxes mark the 25th and 75th percentiles, the whiskers extend to the most extreme nonoutlier data points, and the red lines indicate the medians across mice. The same statistical tests as in (C) are displayed. (**E**) Inactivation effects on trials in which mice initiated movements. The fraction of trials that mice initiated a movement but failed to reach the target, movement onset time, and peak velocity in control versus inactivation trials at the three learning stages. (**F**) Inactivation-induced changes in the fraction of failure, movement onset time, and peak velocity in all initiated movements.

Mice did not move the joystick or mice initiated a movement but failed to reach the target. Mice in the early stage, under M1 inactivation, did not initiate a joystick movement in a larger fraction of trials than control ("Fraction no movements," Fig. 2, C and D). This result alone could indicate that M1 is only important for initiation of a movement. However, in the remainder of trials in which they initiated a movement under M1 inactivation, they failed to move the joystick into the target zone in a majority of the trials ("Fraction failed," Fig. 2, E and F). Furthermore, the onset timing of initiated movements was significantly delayed, and the peak velocity of the movements was also significantly reduced compared to control trials (Fig. 2, E and F). These results indicate that the normal production of the joystick movements in the task relies on M1 during the early stage of training, which is consistent with

previous studies that examined the effects of M1 inactivation on forelimb movements in mice (6, 21, 22).

We also examined the successful movements that entered the target in inactivated trials, although these trials were rare (3.6% of all trials, because inactivation was only in ~12% of trials). We found that the kinematics of these successful movements in inactivated trials were significantly different from successful movements in control trials. Under M1 inactivation, the peak velocity of the movement was reduced, and the length of the path traveled from the origin to the target was elongated compared to control movements (Fig. 3, A and B). These results indicate that, even when mice were able to reach the target under M1 inactivation, the mice were not able to move the joystick as fast or directly to the target as in control trials. Furthermore, the mice often

needed multiple attempts to reach the target, as the initial attempts often failed and, consequently, movement duration from movement onset to target entry increased (Fig. 3, A and B). Therefore, even in successful trials, M1 inactivation in the early learning stage impaired the ability of the mice to produce efficient movements.

In addition to the errors and altered kinematics described above that are evident from the examination of the joystick position, we also noticed from visual inspection that M1 inactivation in the early learning stage often caused mice to lose their grip on the joystick (Fig. 3C). To quantify this effect, we performed video recording of the forelimb
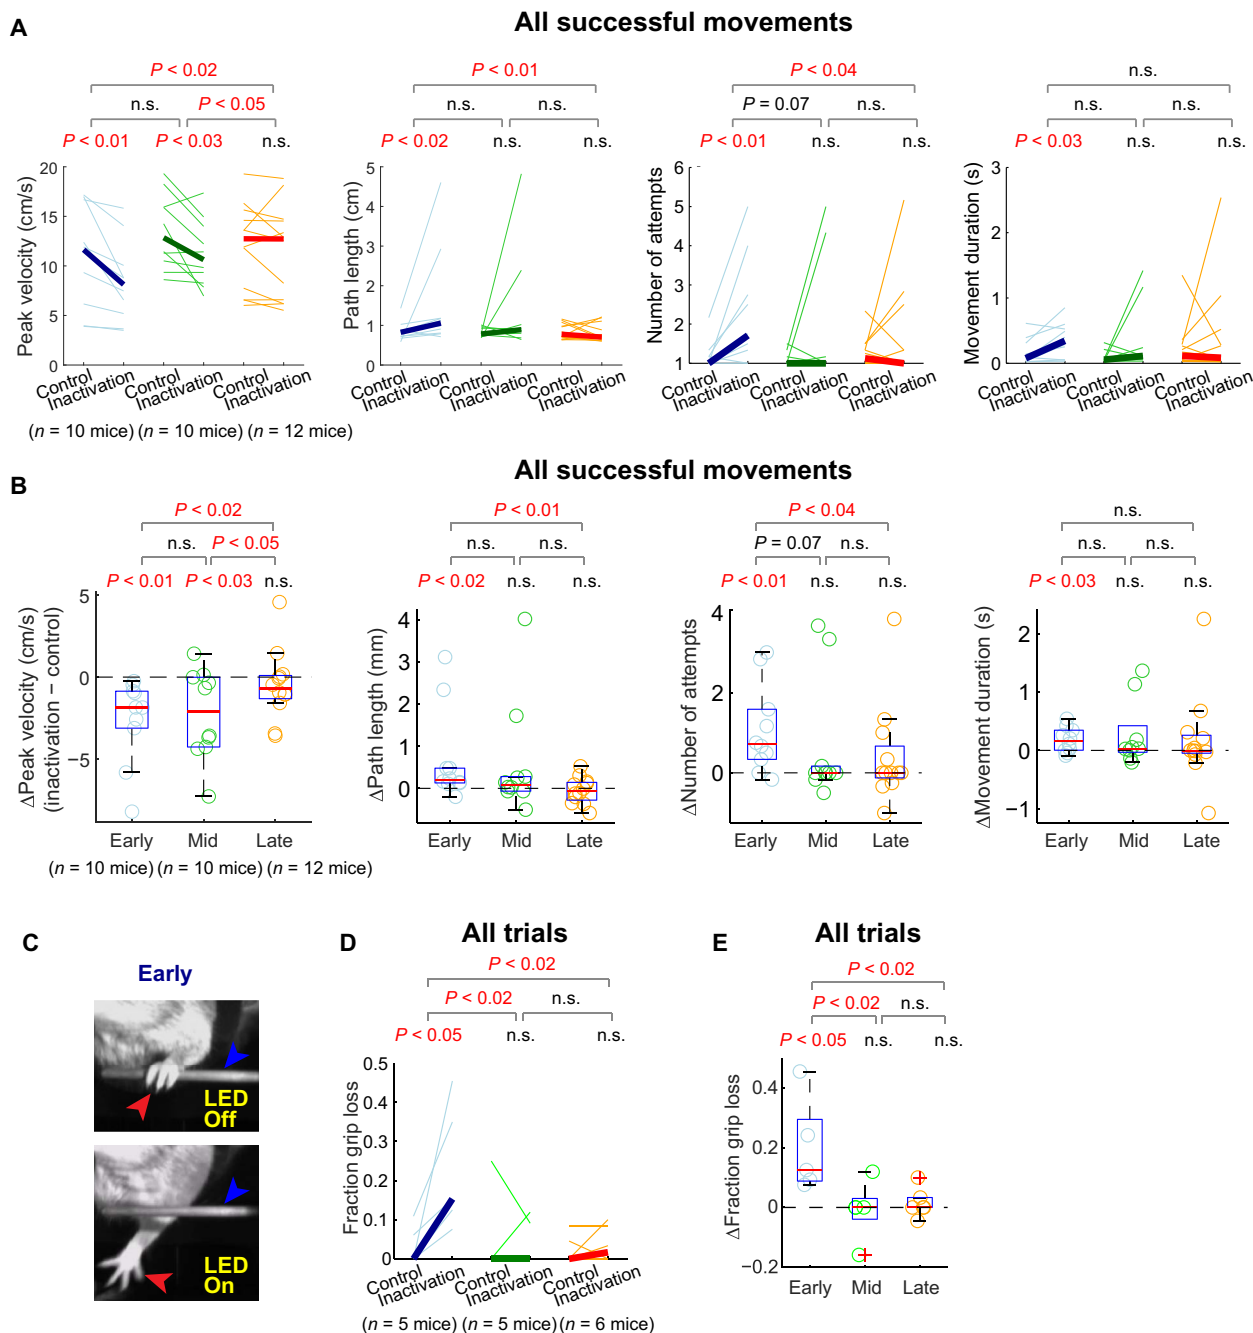
Fig. 3. M1 inactivation affects successful movements and grips on the joystick at the early stage. (**A**) Inactivation effects on movements that successfully entered the target. The peak velocity, path length, number of attempts to reach the target, and movement duration, in control versus inactivation trials, at the three learning stages. Mice that did not make any successful movement under inactivation were excluded from this analysis. Two-sided pairwise comparison between control and inactivation trials within each learning stage are displayed. For the effect size comparison between different learning stages, one-sided unpaired comparison was performed on the differences between control and inactivation trials in each stage. (**B**) Inactivation-induced changes in the peak velocity, path length, number of attempts, and movement duration of all movements that successfully entered the target. The same statistical tests as in (A) are displayed. (**C**) Example frames from videography showing the forepaw (red arrowhead) and the joystick (blue arrowhead) during a control trial (top) and an inactivation trial (bottom) during the early stage of learning. (**D**) The fraction of trials in which mice lost their grip on the joystick, in control versus inactivation trials, at the three learning stages. (**E**) Inactivation-induced changes in the fraction of trials in which mice lost their grip on the joystick.

movements in a subset of mice ($n = 5$ mice) during their task performance. These video recordings were analyzed post hoc by a published method based on deep learning (23) to track the positions of the joystick and the paw. Using the tracked positions, we quantified the fraction of trials in which mice lost their grip on the joystick (Materials and Methods). We found that M1 inactivation in the early learning stage significantly increased the fraction of these trials (Fig. 3, D to E). Together, the variety of measures presented here consistently indicate that the efficient production of movements involving grabbing and moving the joystick during the early stage of training heavily relies on M1 activity.

In stark contrast to the substantial impairment of movements in the early stage, we did not observe any significant impairment associated with M1 inactivation in the late stage of training ($n = 13$ mice). Under late inactivation, mice still initiated movements and successfully entered the target in almost all trials (Fig. 2, C and D). Furthermore, M1 inactivation did not significantly alter the movement onset time, peak velocity, path length, number of attempts, or movement duration (Figs. 2, E and F, and 3, A and B). The analysis of videography also indicated that M1 inactivation did not cause the loss of grip of the joystick ($n = 6$; Fig. 3, D and E). In addition to the lack of significance in these measures, the effect size was significantly different between early and late inactivation in all but one measure. Overall, the movements made under inactivation in the late stage were not distinguishable from the movements in the control trials. Therefore, the movements that initially required M1 for execution became independent of M1 activity after long-term training.

Last, we examined how inactivating M1 affects movements in the middle stage of learning when animals just reached the expert level of performance and movement consistency. M1 inactivation in the middle stage significantly changed the success rate, movement onset time, and movement peak velocity, similar to the early-stage inactivation (Fig. 2, C to F). That is, mice failed to reach the target more frequently, delayed movement initiation, and moved more slowly when M1 was inactivated in the middle stage of learning compared to the control trials at the same stage. However, it is notable that the effect size was significantly reduced in most measures compared to the early stage, such as the fraction of no movements, fraction of failed movements, and peak velocity (Fig. 2, C to F). Movements that successfully entered the target did not show a significant difference between inactivation and control trials in the middle stage, except for the peak velocity (Fig. 3, A and B). Therefore, although M1 inactivation in the middle stage impaired the execution of movements significantly, the effect size was reduced compared to the early-stage inactivation, indicating that M1 dependence of movements decreased from the early to middle stage of learning.

A subset of mice in the middle-stage ($n = 2$) or late-stage ($n = 6$) inactivation groups were also used in the early stage inactivation experiment, raising a possibility that the smaller inactivation effects in later stages might be mediated by these mice, via some compensatory mechanisms acquired from their earlier experience of inactivation. However, this seems unlikely because we observed decreasing inactivation effects even in mice without a prior experience of inactivation (fig. S4). Together, we found that M1 inactivation effects on movements gradually decrease from the early to middle and to late learning stages to the extent that movements after long-term training can be produced normally even when M1 is inactivated.

M1 activity exhibits biphasic evolution during long-term learning

To examine how M1 activity may change while its involvement in movements declines with learning, we performed longitudinal two-photon calcium imaging during the joystick task over the course of 60-day training ($n = 5$; Fig. 4A). Similar to the larger set of mice shown in Fig. 1, task performance and kinematic stereotypy of the imaged mice also exhibited improvement with training (Fig. 4B). The imaged mice reached their expert levels in 24 ± 5 days (mean ± SD), similar to the larger set (23 ± 4 days). Accordingly, behaviors after day 19 were significantly better than the beginning of training (days 1 to 3) but similar to the end of training (days 50 to 60) in all but one measure (Fig. 4B).

Using two-photon calcium imaging, we recorded the activity of layer 2/3 excitatory neurons in the right M1 during the joystick task (Materials and Methods). We repeatedly imaged the same field of neurons every day (Fig. 4C). However, to maximize yield, we analyzed all neurons in the field each day regardless of whether they could be consistently identified across different imaging days. The transgenic expression of GCaMP6s (*CamkIIa-tTA::tetO-GCaMP6s*) allowed stable longitudinal imaging (Fig. 4C). Here, we sought to compare M1 population activity across the three learning stages at which we examined the effects of M1 inactivation in the earlier section. Similar to the time periods of the inactivation experiments, we defined days 4 to 8 as the early, days 20 to 25 as the middle, and days 50 to 60 as the late stage. During the middle stage, mice reached their expert levels as shown above.

We first examined the signal-to-noise ratio (SNR) in each imaging day to check whether the recording quality is stable over time (Materials and Methods). We found no significant changes in SNR, supporting the stability of imaging. The number of analyzed neurons was also not significantly different across the three stages (Fig. 4D). Nevertheless, to avoid potential sample size effects when comparing population activity across different learning stages, we matched the number of neurons across different populations by comparing their subpopulations with 50 randomly selected neurons (Materials and Methods).

It has been shown that M1 population activity associated with learned movements becomes gradually more consistent across trials over 2 weeks of training in a task with a one-dimensional lever but otherwise similar to the current joystick task (6). That is, the same set of neurons is more reliably recruited for the production of learned movements at a later stage compared to the early stage. This was interpreted as an emergence of neural ensembles in M1 dedicated to the production of the learned movement. To examine this learning-associated change, we computed the correlation coefficient of the population activity for every pair of successful movements in each training day, following the previous method (Fig. 4E). In line with the previous report, we observed that the trial-to-trial correlation of M1 population activity significantly increased in the early phase of training during which task performance and kinematic stereotypy rapidly improved (comparison between the early and middle stages; Fig. 4F). Unexpectedly, however, the improvement in activity consistency during the early phase of learning was not maintained with prolonged training but instead was followed by a decrease (comparison between the middle and late stages; Fig. 4F). The reduction of activity consistency occurred despite the sustained expert motor behaviors during these two expert stages (Fig. 4B).

To more closely examine the relationship between motor behaviors and M1 population activity, we analyzed the relationship between the correlation of population activity and the correlation of movement trajectories for pairs of successful trials, throughout the course of 60-day training, following a previous method (Fig. 5A) (6). We observed a general pattern that more similar movements (i.e., more correlated movements) are associated with more similar population activity (i.e., more correlated activity) in all learning stages. However, the similarity
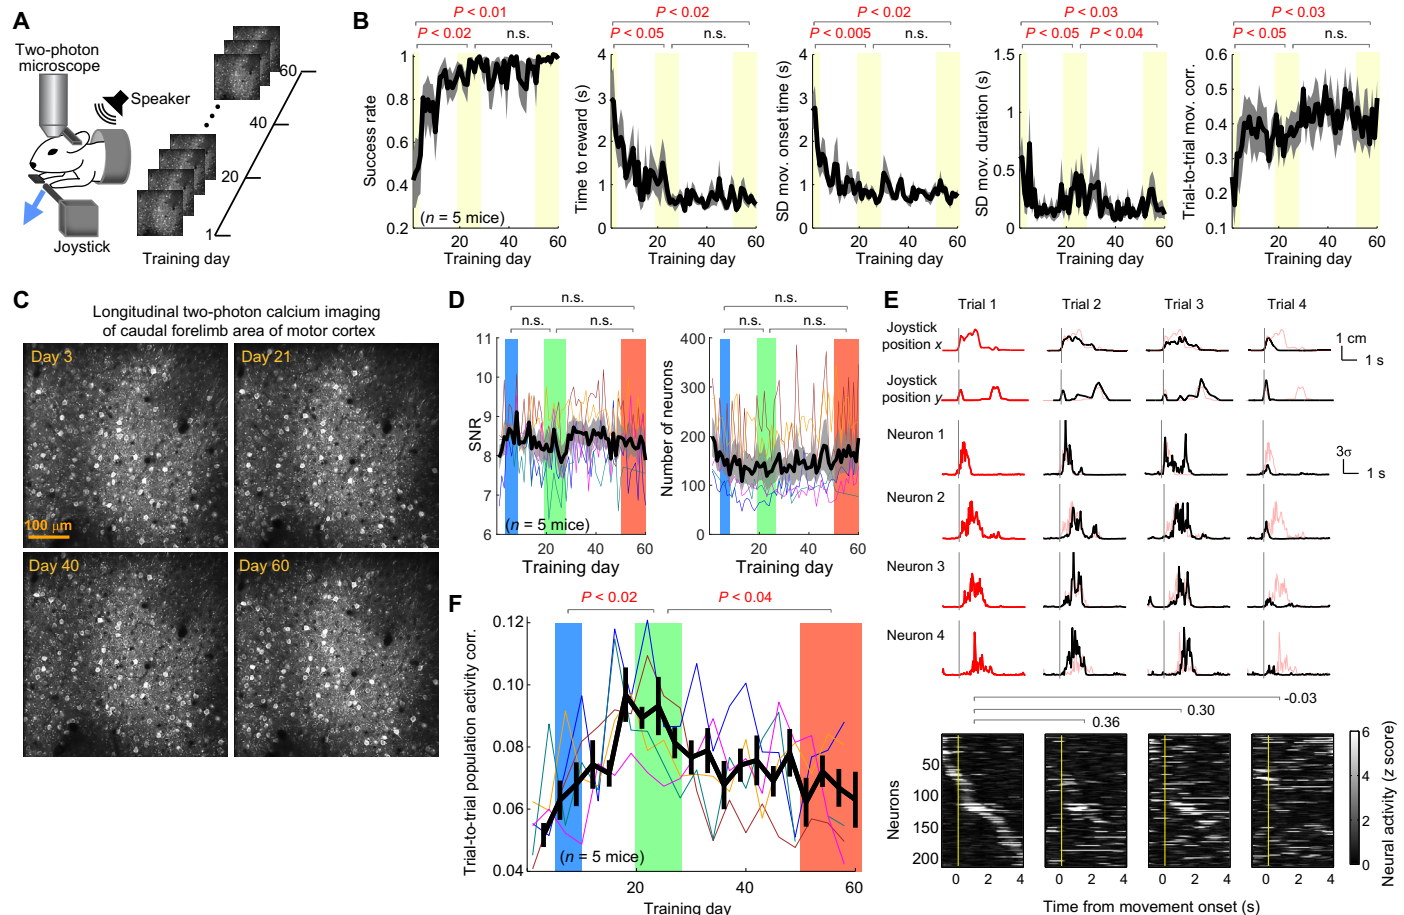
Fig. 4. M1 population activity consistency evolves in two phases over long-term training. (**A**) Longitudinal imaging of the neurons in the same field in M1 over the course of 60-day training in the joystick task. (**B**) Task performance and movement consistency during the 60-day training in the imaging mice: the success rate, time required to acquire reward, variability of movement onset time, variability of movement duration, and trial-by-trial movement trajectory correlation as a function of training day, from left to right. Mean ± SEM ($n = 5$ mice). The three shaded regions correspond to days 1 to 3, 19 to 29, and 50 to 60, respectively. Two-sided pairwise comparison for each pair of learning stages is displayed. (**C**) The imaging fields from training days 3, 21, 40, and 60 of a single mouse. (**D**) The SNR and the number of neurons in the imaging field as a function of training day. Mean ± SEM ($n = 5$ mice). The blue, green, and red shaded regions correspond to the early, middle, and late learning stages in the inactivation experiments, respectively. (**E**) Movement trajectory, four single-neuron activity, and whole-population activity (heat map) from four example trials in a single day. The vertical lines mark movement onset. The trajectory and single neuron activity of trial 1 were superposed in the other trials as a thin red line for comparison. The neurons in the heat map are sorted in the same way in all four trials according to the peak activity time in trial 1. The three numbers above the heat maps indicate the trial-to-trial population activity correlation for the corresponding trial pairs. (**F**) The trial-to-trial population activity correlation as a function of training day (bin size: 3 days). Thin different colors represent different mice, and the black line represents mean ± SEM ($n = 5$ mice). The three shaded regions include bins corresponding to the three learning stages in (D). Two-sided pairwise comparison for each pair of learning stages is displayed.

of population activity for a given movement similarity was not constant across learning stages (Fig. 5B). Instead, it increased from the early to middle stage and then decreased from the middle to late stage, at all levels of movement similarity (Fig. 5B). In particular, activity consistency for highly similar movements indicates that the relationship between movements and M1 activity shows the least degeneracy during the middle learning stage.

The decreased consistency and increased degeneracy from the middle to late stage occurred despite the maintained expert performance of joystick movements. To test whether other movements became more variable during these expert stages, which could contribute to the increased variability of M1 activity, we analyzed the licking patterns. Contrary to this possibility, however, we found that licking patterns in our task became more consistent during the prolonged training in the expert stages (fig. S5). Thus, although we cannot exclude the possibility

that some unmeasured movements became more variable during the middle to late learning stages, we favor the interpretation that the increased variability of M1 population activity reflects the dissociation of M1 activity from forelimb movement control.

Last, we examined how the fraction and activity level of movement-related neurons changed over time (fig. S6). Movement-related neurons were defined as those with significantly different activity between movement and baseline periods (Materials and Methods). We found that the fraction of movement-related neurons increased in the early phase of learning (days 1 to 21, Pearson's correlation coefficient, $r = 0.23$, $P < 0.03$) and then decreased in the later phase (days 22 to 60, $r = -0.21$, $P < 0.01$), echoing the biphasic pattern observed in the population activity consistency. We also examined the activity level of the movement-related neurons during each movement. A substantial amount of day-to-day fluctuation was apparent in the movement period
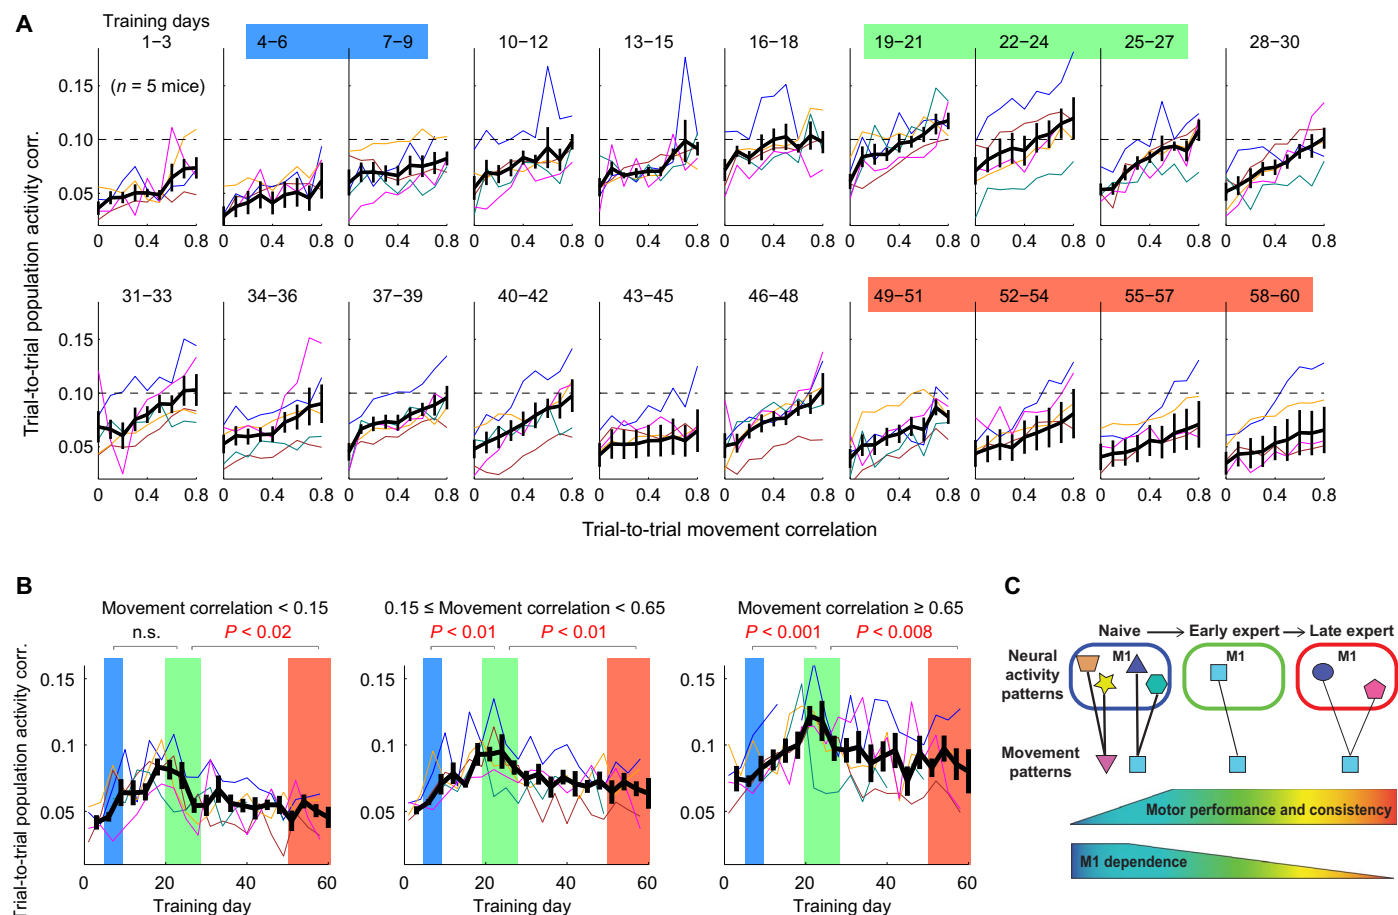
Fig. 5. The relationship between M1 activity and movements varies with learning stages. (**A**) The relationship between trial-to-trial movement trajectory correlation and population activity correlation over the 60-day training (bin size: 3 days). The x axis represents the movement correlation binned in nine intervals, and the y axis represents the mean population activity in each interval. Thin different colors represent different mice, and the black line represents mean ± SEM ($n = 5$ mice). (**B**) The evolution of population activity consistency over 60 days for movement pairs with low, middle, and high movement trajectory correlations, from left to right. The three shaded regions are the same as the three learning stages in Fig. 4F. The increase followed by decrease of activity correlation is observed consistently in all groups of movement correlations. (**C**) A model for the evolution of M1 engagement over the course of long-term motor learning. The thickness of the lines between neural activity patterns and movement patterns indicate the degree of M1 dependence of movements inferred from the inactivation experiments. The consistency and degeneracy between activity and movement patterns are derived from the imaging experiments.

activity. Nevertheless, we found a general pattern that the activity of movement-related neurons gradually decreased over the course of the 60-day training (days 1 to 60, $r = -0.16$, $P < 0.01$).

Together, we found that the early phase of learning accompanied an increased consistency and decreased degeneracy of M1 population activity and an increased fraction of movement-related neurons. However, these changes associated with the early-phase learning were not maintained during the later phase of training, despite maintained motor behaviors. Furthermore, the average activity of movement-related neurons gradually decreased. These changes in M1 activity and the decaying M1 dependence of movements support the notion that the involvement of M1 in movement control dynamically varies with learning stages.

DISCUSSION

In this study, we found that, in a forelimb motor learning task, M1 is essential for movement control during the early phase of learning, but M1 becomes gradually disengaged from movement control over

months of training. These results indicate that the brain contains multiple movement control systems, one involving M1 and another by-passing M1, and the system bypassing M1 can increasingly take over the control of movements as learning progresses (Fig. 5C).

Our longitudinal imaging experiment revealed that M1 population activity showed biphasic changes, while movements became gradually independent of M1. The early phase involved an improved consistency of population activity and decreased degeneracy in the relationship between M1 activity and movements, but these changes were gradually lost during the later phase, despite the maintained expert-level performance. Several factors could contribute to the improved consistency of population activity during the early phase. First, movements became significantly more consistent during this time period (Fig. 1). Second, even for equally similar movements, associated population activity became more consistent, i.e., reduced degeneracy in the relationship between movements and M1 activity (Fig. 5). Third, well-prepared movements have been shown to accompany less variable activity in M1, raising a possibility that mice learned to better prepare movements in our task during this early phase (24, 25).

The substantial changes in M1 activity during long-term training in the expert stage indicate that, even when there are little overt changes in behavior, the underlying neural control system could be changing. Conversely, although the level of activity consistency in the early and late stages is similarly lower than in the middle stage, motor behaviors greatly differ between the two stages. Our inactivation experiment results suggest that the underlying source of the low activity consistency may differ between the early and late stages. The low consistency during the early stage might reflect a high variability of movements and an intrinsically redundant and degenerate relationship between M1 activity and movements, manifesting during an exploration in the activity space when searching activity patterns that generate a desirable movement. Once an activity pattern that generates a desirable movement is found, this activity pattern might be reinforced over time to be consistently recruited during the task performance. We observed increased consistency in M1 activity from the early to middle stage. In contrast, during the further training in experts, movements become less dependent on M1, and thus M1 activity would be dissociated from movements, and such dissociation would permit the consistency to decrease (Fig. 5C).

The effects of M1 inactivation seemed to decline gradually throughout the long-term training, and the middle-stage inactivation showed smaller effects compared to the early stage in some of the measures (Figs. 2 and 3). Somewhat counterintuitively, the gradual decrease in M1 dependence occurs concurrently with the learning-related refinement of M1 activity during the early phase of learning. Thus, we propose that the early phase of learning involves two parallel changes: First, M1 activity becomes more refined to reliably drive the learned movement; and second, an alternative pathway bypassing M1 becomes gradually entrained, which could eventually take over M1's role in movement generation (Fig. 5C).

The transition from M1-dependent to M1-independent movement control during long-term learning might be related to the automatization of movement execution, a well-known phenomenon following long-term training. In early stages of learning, a considerable amount of cognitive effort is allocated to finding apt strategies that efficiently achieve the goal and making adjustments (26). However, at the expert stage, highly skilled movements can be automatically generated with little conscious effort, resembling innate behaviors (27). Such a shift can be readily appreciable in many motor skills supporting our daily activity, such as typing on a computer keyboard. At the highly trained stage, movement execution becomes so automatic that conscious attention can even disrupt the execution of the overlearned skills, causing performance decrements in some cases (28). It has been suggested that effortless movement control is tied to a reduced level of engagement of M1 and other frontal cortical areas (27, 29, 30). Our finding that M1 engagement decreases with long-term training may reflect a transition toward a less effortful, more automatic execution of the highly learned movement.

Our observation that M1 is required for movement control early in learning does not distinguish whether M1 is causally driving the movement or is simply permissive for the rest of the circuit to generate the movement. Furthermore, our results in no way indicate that any movement can become M1 independent with long-term training. Certain movements will always remain under M1 control even after years of training, while certain movements never require M1. In addition, we should use caution when extending the current results to other species. M1 lesions tend to cause severe and long-lasting movement deficits in primates, especially humans (31, 32), while the effect of M1 lesions in rodents is more nuanced (33). Therefore, M1 functions and the degree of M1 dependence of movements likely differ across species, and it has yet to be tested whether some M1-dependent movements in primates can also become M1 independent with long-term learning. However, the partial recovery seen with rehabilitation after M1 insults in humans suggests that at least some M1-dependent movements can become M1 independent.

Our finding that the production of movement is dependent on M1 early but not late in training is distinct from a recent report that M1 is not required for executing movements at any time throughout learning (33). Kawai et al. (33) instead showed that M1 is only required for learning, or the improvement, of a motor skill. The difference in these studies may be related to the difference in the movements studied (originally M1-dependent versus originally M1-independent movements) and/or the perturbation methods (optogenetic inactivation versus lesion). Nevertheless, our observation that the same inactivation leads to markedly different behavioral effects indicates that long-term learning can offload the function of movement control from M1. The brain structures that control the movements in expert animals may include the brainstem, cerebellum, and basal ganglia, all of which contain descending pathways to control motor circuits in the spinal cord (34, 35). In addition to its role in executing movements early in learning, M1 may also play an instructive role in the offloading process by acting as a tutor for the subcortical structures (33). It has been proposed that behaviorally relevant circuits could have an additional function to entrain shortcut circuits through Hebbian plasticity (36, 37). For instance, the consistent M1 activity patterns during the middle stage of learning might be critical for training the alternative, shortcut circuits.

The potential advantage of having multiple movement control systems for different stages of learning is unknown, but we favor a hypothesis that each system has a distinct level of stability and flexibility. M1, with its high levels of plasticity, may be best suited to encode newly acquired skills, which may need to be modified during initial stages of learning. When the skill is highly learned and the need for further modification is reduced, it may be beneficial to offload the movement to a more stable system, which may allow automatic and reproducible movements and also make M1 available to learn other novel movements. This sequential learning process fits well with the two-phase changes that we found in M1 activity and resembles other systems of learning. For example, episodic memory is believed to be initially stored in hippocampus and later transferred to cortex (37). Furthermore, reward-guided operant learning initially depends on the dorsomedial striatum, but later, the dependence shifts to dorsolateral striatum as the behavior becomes habitual (16). In both contexts, the dynamic areas initially important for learning may offload the information to other areas for stable long-term storage. Our results in M1 suggest that the dynamic and fluid nature of learning circuits might be a fundamental scheme for long-term learning.

MATERIALS AND METHODS
Animals
All procedures were in accordance with protocols approved by the University of California, San Diego Institutional Animal Care and Use Committee and the guidelines of the National Institutes of Health. Mice (6 weeks or older, male and female, calcium imaging: cross between *CaMK2a-tTA* [JAX 003010] and *tetO-GCaMP6s* [JAX 024742]; optogenetic inactivation: cross between *PV-Cre* [JAX 00869] and *Ai32* [JAX 024109]) were housed in a room with a reversed light cycle

(12 hours–12 hours). Experiments were performed during the dark period.

Behavioral apparatus

The behavioral apparatus was housed in a soundproof box (40 cm by 40 cm by 40 cm), and the joystick task was performed in the dark. The components of the task (17) included a joystick (M11L061P, CH Products) and a water port (with photodiodes to sense licking). The joystick handle was custom machined and fitted with a 1.6-mm-thick brass rod that mice manipulated with their left forepaw. An electromagnet (EM050-3-222, APW) mechanically immobilized the joystick at the origin during intertrial intervals. The joystick had a dynamic range of 5 cm in each of two directions. The two-dimensional position of the joystick was continuously recorded at 1 kHz using a data acquisition card (USB6008, National Instruments) and custom MATLAB software. The task sequence execution, auditory cue presentation, and reward dispensation were coordinated (and recorded) by an open source real-time Linux/MATLAB software package BControl (http://brodywiki.princeton.edu/bcontrol/).

Behavioral training of the joystick task

In the joystick task, the joystick was released from the electromagnet immobilization at the beginning of each trial. Two seconds after the trial onset, a 6-kHz auditory tone was played. If mice moved the joystick into the target within 10 s from the auditory tone onset, then they received a reward, even in trials where they initiated movements before the auditory tone. The return of the joystick to the origin ended the trial and initiated an intertrial interval (4 s), during which the joystick was immobilized at the origin by the electromagnet.

Before mice started the training for the task, they were familiarized with an easier version of the task with a larger single target zone covering the whole angular range of the joystick. Thus, displacement of the joystick from the origin by approximately 6 mm in any direction was considered a target entry. The mice were trained in the easy task until they acquired reward in at least 70 trials of 100. This criterion was reached in 2 to 7 days.

The main task was identical to the familiarization task except that the target zone was reduced to cover only 80% of the joystick's dynamic range, excluding each edge area. Therefore, mice could not ride edges all the way to reach the target. Taking a further cautious step in our analysis, we excluded any trials during which movement was along an edge for more than half the target distance. The number of trials meeting the analysis inclusion criterion did not significantly change with training. All mice were presented with 100 trials per day for 60 days, except for mice that were trained only for the early or middle stage inactivation experiments.

Movement analysis

Joystick position–related events and kinematic variables were defined and measured as described below.

Movement onset: The first time at which the joystick velocity exceeded 20 mm/s continuously for 20 ms and the joystick moved at least 1.1 mm from the origin. Movement onset time: Time from the trial onset to movement onset. Target entry: The first time at which the joystick enters the target zone since the most recent departure from the origin. Target entry could occur more than once in a single trial. Target entry in the main text refers to the first target entry unless otherwise noted. Movement duration: Time from movement onset to the first target entry. Movement offset: The first time when the joystick velocity fell below

20 mm/s continuously for 20 ms since the final target entry. Trial-by-trial movement correlation: Correlation coefficient between the two joystick traces (the concatenated x and y position time series from −1 to 4 s from movement onset). The time window for the trajectory correlation analysis was chosen to cover the period from movement onset to movement offset for over 90% of all successful trials (90th percentile, 3.6 ms; 92.5th percentile, 4.0 ms). This time period may include return movements back to the start position after target entry. Path length: Velocity integrated from movement onset to the first target entry. Number of attempts: The number of peaks in the velocity trace from the movement onset to the first target entry.

Licking behavior

We recorded the licking behaviors of mice during the task using a custom infrared beam–based sensor placed in front of the lick port. By detecting the times at which the infrared beam was interrupted by a tongue protrusion, we created a time series of lick events. The lick event time series were aligned to the forelimb movement onset in each trial, spanning the same time period of M1 activity analysis (−1 to 4 s from movement onset). The similarity of lick patterns between trials was measured by computing the correlation coefficient between the movement onset aligned lick event time series.

Inactivation experiment

Mice (PV-Cre::Ai32) used for inactivation experiments were implanted with head-bar and cranial windows over the forelimb region of M1 bilaterally (coordinates relative to bregma: ±1.5 mm lateral, +0.3 mm anterior). Following a minimum 3 days of recovery, daily water consumption was limited to a controlled volume (typically 1 ml/day). After 3 to 10 days of water restriction, the mice began behavioral training.

For the early-stage inactivation experiment ($n = 13$ mice), 3 days were randomly selected between days 4 and 8 for M1 inactivation and 3 other days between days 1 and 9 for head-bar control (Fig. 2, A and B). For the middle-stage inactivation ($n = 10$ mice), 3 days were randomly selected between days 20 and 25 and 3 other days between days 19 and 26 for control. For expert-stage inactivation ($n = 13$ mice), 3 days were randomly selected between days 61 and 69 and 3 others for control. Six mice were used for both early and late inactivation. Two mice were used for both early and middle inactivation. Five mice were used only for early-inactivation experiment. Eight mice were used only for middle-inactivation experiment. Seven mice participated only in the late-inactivation experiment. The mice were randomly assigned to early-, middle-, or late-inactivation groups. The cranial windows were cleaned with cotton swabs and ethanol and visually inspected for their clarity before each inactivation experiment began. In all mice, blood vessels and dura underneath the windows were visible with the naked eye.

In M1 inactivation sessions, the distal ends of a bifurcated patch cord (Doric Lenses) were placed directly on the cranial windows, and blue LED light (465 nm, ~3.75 mW at each end, LEDC1-B_FC and LEDRV_1CH_1000, Doric Lenses) was delivered on a randomly selected 12% of trials. Head-bar day experiments were identical to the inactivation days except that the patch cord ends were placed ~1 mm above the head bar, away from the cranial windows. To control for any nonspecific light effects, we used light-on trials on the head-bar days as control trials in all our analyses.

Longitudinal two-photon calcium imaging experiment

Mice (CaMK2-tTA::tetO-GCaMP6s) used for imaging experiments were implanted with a head plate and a cranial window over the forelimb

region of M1 on their right hemisphere, and then underwent the recovery and water restriction procedures described above. After 2 to 7 days of task familiarization as described in the section "Behavioral training of the joystick task," we started imaging cortical activity with excitation at 925 nm from a Ti-Sa laser (Spectra-Physics) at ~28 frames/s using a two-photon microscope (B-SCOPE, Thorlabs). For each mouse, a single field of view in the forelimb region of M1 (covering 472 μm by 508 μm at a depth of approximately 250 μm beneath the dura in layer 2/3) was longitudinally imaged over the course of 60-day training. Although a single field of view was imaged throughout the experiment, data from each day were processed independently without limiting our analyses to neurons present in all days. Only the imaging days with satisfying image clarity and no other technical issues were analyzed (54 ± 3 days, mean ± SD across five mice).

Single-cell activity

Using a custom MATLAB program, fluorescence images were aligned frame by frame to compensate for lateral motions post hoc (38). Regions of interest (ROIs) were manually drawn on the motion-corrected fluorescence images, by circumscribing the cell bodies based on their GCaMP fluorescence intensity distinguishable from the background. Pixels inside each ROI were considered as a single soma, whereas pixels extending radially outward from the cell boundary by 2 to 6 pixels were considered background. For each ROI, we subtracted 70% of the average background pixel intensity from the average soma pixel intensity at each frame as the fluorescence signal of the ROI. The fluorescence signals were transformed to dF/F following the procedure in the previous study (17) and then further transformed into an estimate of spike rates using the spike-triggered mixture model (https://github.com/lucastheis/c2s) (39).

Signal-to-noise ratio

For each ROI, we first computed the mean (μ) and SD (σ) of its fluorescence signal, and detected all calcium events using MATLAB function findpeaks. In this function, the minimum peak height was set to be μ + 2σ, and the minimum distance between adjacent peaks was set to be 10 frames (~350 ms). Surrounding each detected peak, we delimited its event period as the time period in which the fluorescence signal was continuously above μ. We treated the signal outside the event periods as noise. Using the detected peak heights and noise, we computed SNR for each ROI as the following: $\mathrm{SNR} = \frac{\mathrm{mean(peak\ heights)} - \mathrm{mean(noise)}}{\mathrm{SD\ (noise)}}$.

Trial-to-trial population activity correlation

The population activity correlation between two trials was the correlation coefficient between the two concatenated activity time series (−1 to 4 s from movement onset) of a population of neurons (Fig. 3E). Since the number of neurons in a population varies across days and animals (Fig. 4D), we matched the population size by randomly subsampling 50 neurons and computed the trial-to-trial correlation of the 50-neuron population activity. The subsampling process was repeated 100 times, and the average across 100 trial-to-trial correlations was used for the given population.

Relationship between movements and population activity

For each pair of trials, we computed the correlation coefficients between the movement trajectories and between the population activity (population size matched as described above) in each day. Pairs of trials in a 3-day bin were pooled together, and 1000 pairs were randomly sampled. The 1000 pairs were binned into nine intervals based on movement correlation, two boundary and seven intermediate intervals between −1 and 1 (Fig. 5A). The lower boundary interval included all the pairs in which movement correlation was less than 0.05, the upper boundary interval included all the pairs with movement correlation greater than

or equal to 0.75, and the intermediate intervals were uniformly spaced between 0.05 and 0.75.

Movement-related neurons

For each neuron, we tested whether the distribution of its movement period activity was significantly different from that of baseline activity using Wilcoxon rank sum test ($P < 0.01$). We defined baseline as the 0.5-s period before trial onset. A wide range of movement periods were examined, ranging from 0.5- to 4-s windows after movement onset. A similar and significant trend was seen across all different time windows, except that a decrease in the fraction during the later phase was statistically significant up to 1 s. Data presented in fig. S6 are from the 0.5-s window.

Extracellular electrophysiology

Extracellular recordings were performed similar to those previously described (40). Adult mice (PV-Cre::Ai32, $n = 2$), 6 weeks or older were anesthetized with urethane (1.2 g/kg, intraperitoneal) and given the sedative chlorprothixene (0.05 ml of 4 mg/ml, intramuscular) and implanted with a T-shaped head bar for head fixation. Body temperature was maintained at 37°C using a feedback-controlled heating pad (40-90-8D, FHC Inc.). A uniform layer of silicone oil was applied to the eyes to prevent drying. A craniotomy ~1 mm in diameter was made over the middle of V1 (~2.75 mm lateral to the midline and ~0 mm anterior to the lambda suture), and sterile saline was placed in the well of the craniotomy to keep the brain moist. A 16-channel linear silicon probe (a1x16-5 mm-25-177, NeuroNexus) mounted on a manipulator (Luigs & Neumann) was slowly advanced into the brain to a depth of ~750 μm. Recordings were started 20 min after insertion of the probe into V1. Signals were amplified 400-fold, band-pass–filtered (0.3 to 5000 Hz, with the presence of a 60-Hz notch filter, A-M Systems 3600), and then digitized at 32 kHz (PCIe-6259, National Instruments) with custom MATLAB software.

Visual stimulus was presented across three computer monitors (VX2450wm-LED, 60-Hz refresh rate, gamma corrected, ViewSonic) mounted orthogonally to each other to form a square enclosure that covered ~270° of the visual field along the azimuth. The mouse head was immobilized at the center of the enclosure. Visual stimuli were generated using Psychtoolbox. The gratings drifted clockwise or counterclockwise in an oscillatory manner (amplitude ± 5°; grating spatial frequency, 0.08 cycles per degree; oscillation frequency, 0.4 Hz; contrast, 100%; mean luminance, 40 cd/m^2). Trials were spaced by an interstimulation interval of 8 s.

Optogenetic stimulation of V1 was accomplished by shining 470-nm blue light through an optical fiber pointed at V1. We recorded from V1 using three different blue light intensities: 3.5, 7.0, and 10.5 mW. Blue light intensities were varied in separate blocks of trials (i.e., 100 trials of 3.5 mW, followed by 100 trials of 7.0 mW). During optogenetic cortical inactivating trials, 10 s of blue-light stimulation were applied in the middle of 12 s of visual stimulus. Trials of cortical inactivation (light on) were interleaved with control trials (light off).

Multiunit analysis

Multiunit activity was isolated using spike-sorting software in MATLAB as previously described (40). The raw extracellular signal was band-pass–filtered between 0.5 and 10 kHz. Spiking events were detected with a threshold of 3.5 times the SD of the filtered signal. Spike waveforms of four adjacent electrode sites were clustered using a k-means algorithm. Multiunit spiking activity was defined as all spiking events exceeding the detection threshold after the removal of electrical noise

or movement artifacts by the sorting algorithm. Individual spiking events were assigned to one of the 16 recording sites according to where they showed the largest amplitude.

Video analysis of forelimb movements

Mice performing the task were video recorded at the rate of 30 frames/s, with the resolution of ~0.15 mm per pixel (DMK 23U618, Imaging Source). Five points of interest (POIs) that we tracked in each frame were the tip positions of the three digits on the radial side (analogous to index, middle, and ring fingers) of the left paw and the two end points of the linear joystick bar. Lighting conditions and camera angles slightly differed across days, so we manually sorted all recording days into five groups, each with similar recording settings, and built a deep neural network model separately for each group. In each group, we first randomly selected 180 frames and manually labeled POIs in those frames and used them to train and test a neural network model implemented in DeepLabCut (23). The trained model tracked the POIs in the test data with an average tracking error of less than 2.5 pixels. Applying the trained model to unlabeled frames produced the POIs and the strength of evidence (range, 0 to 100) for each POI in each frame. Then, in each frame, we identified the closest digit from the line between the two end points of the joystick and deemed the frame as high confidence if the strength of evidence for the identified digit was greater than 74. On the basis of this criterion, 94% of frames were classified as high confidence. In each high-confidence frame, we calculated the distance between the closest digit and the joystick bar. The distances in low-confidence frames were linearly interpolated using the nearest high-confidence frames. On a given trial, we declared a grip loss if the distance was greater than 20 pixels (~3 mm) for at least 15 consecutive frames (0.5 s). In very rare trials that included more than 10 consecutive low-confidence frames (~ 2%), grip losses were manually scored.

Statistical analysis

For within-condition comparisons, we applied either a Wilcoxon signed-rank test or Student t test on a set of paired values from each animal (e.g., control versus inactivation trials in early-stage inactivation), depending on the result of Lilliefors goodness-of-fit test with the null hypothesis that the data were normally distributed ($P < 0.05$ was used for the rejection of the null hypothesis). For effect size comparisons between conditions (early-stage inactivation versus late-stage inactivation), we applied a Wilcoxon rank sum test or two-sample t test as the two samples were not from identical sets of animals.

SUPPLEMENTARY MATERIALS

Supplementary material for this article is available at http://advances.sciencemag.org/cgi/content/full/5/10/eaay0001/DC1

Fig. S1. Reaching movements are goal directed after 60 days of training.
Fig. S2. Optogenetic activation of PV inhibitory neurons inactivates cortex.
Fig. S3. Task performance and movements in light-off trials in head-bar control versus M1 inactivation days are generally equivalent.
Fig. S4. M1 inactivation effects are reduced in the later learning stages compared to the early stage, even in mice without a prior experience of early-stage inactivation.
Fig. S5. Licking variability does not explain the longitudinal changes in M1 activity consistency.
Fig. S6. The fraction of movement-related neurons and their activity level also change over the course of long-term learning.

View/request a protocol for this paper from Bio-protocol.

REFERENCES AND NOTES

1. J. N. Sanes, J. P. Donoghue, Plasticity and primary motor cortex. Annu. Rev. Neurosci. 23, 393–415 (2000).

2. A. J. Peters, H. Liu, T. Komiyama, Learning in the Rodent Motor Cortex. Annu. Rev. Neurosci. 40, 77–97 (2017).

3. T. Xu, X. Yu, A. J. Perlik, W. F. Tobin, J. A. Zweig, K. Tennant, T. Jones, Y. Zuo, Rapid formation and selective stabilization of synapses for enduring motor memories. Nature 462, 915–919 (2009).

4. G. Yang, F. Pan, W.-B. Gan, Stably maintained dendritic spines are associated with lifelong memories. Nature 462, 920–924 (2009).

5. M. Fu, X. Yu, J. Lu, Y. Zuo, Repetitive motor learning induces coordinated formation of clustered dendritic spines in vivo. Nature 483, 92–95 (2012).

6. A. J. Peters, S. X. Chen, T. Komiyama, Emergence of reproducible spatiotemporal activity during motor learning. Nature 510, 263–267 (2014).

7. S. X. Chen, A. N. Kim, A. J. Peters, T. Komiyama, Subtype-specific plasticity of inhibitory circuits in motor cortex during motor learning. Nat. Neurosci. 18, 1109–1115 (2015).

8. M. Rioult-Pedotti, D. Friedman, J. Donoghue, Learning-induced LTP in neocortex. Science 290, 533–536 (2000).

9. R. M. Costa, D. Cohen, M. A. L. Nicolelis, Differential corticostriatal plasticity during fast and slow motor skill learning in mice. Curr. Biol. 14, 1124–1134 (2004).

10. D. Huber, D. A. Gutnisky, S. Peron, D. H. O'Connor, J. S. Wiegert, L. Tian, T. G. Oertner, L. L. Looger, K. Svoboda, Multiple dynamic representations in the motor cortex during sensorimotor learning. Nature 484, 473–478 (2012).

11. Y. Masamizu, Y. R. Tanaka, Y. H. Tanaka, R. Hira, F. Ohkubo, K. Kitamura, Y. Isomura, T. Okada, M. Matsuzaki, Two distinct layer-specific dynamics of cortical ensembles during learning of a motor task. Nat. Neurosci. 17, 987–994 (2014).

12. A. J. Peters, J. Lee, N. G. Hedrick, K. O'Neil, T. Komiyama, Reorganization of corticospinal output during motor learning. Nat. Neurosci. 20, 1133–1141 (2017).

13. G. Wulf, Attention and Motor Skill Learning (Human Kinetics, 2007).

14. E. Dayan, L. G. Cohen, Neuroplasticity subserving motor skill learning. Neuron 72, 443–454 (2011).

15. A. K. Dhawale, M. A. Smith, B. P. Ölveczky, The Role of variability in motor learning. Annu. Rev. Neurosci. 40, 479–498 (2017).

16. H. H. Yin, S. P. Mulcare, M. R. F. Hilário, E. Clouse, T. Holloway, M. I. Davis, A. C. Hansson, D. M. Lovinger, R. M. Costa, Dynamic reorganization of striatal circuits during the acquisition and consolidation of a skill. Nat. Neurosci. 12, 333–341 (2009).

17. E. J. Hwang, J. E. Dahlen, M. Mukundan, T. Komiyama, History-based action selection bias in posterior parietal cortex. Nat. Commun. 8, 1242 (2017).

18. H. Makino, C. Ren, H. Liu, A. N. Kim, N. Kondapaneni, X. Liu, D. Kuzum, T. Komiyama, Transformation of cortex-wide emergent properties during motor learning. Neuron 94, 880–890.e8 (2017).

19. M. Brus-Ramer, J. B. Carmel, J. H. Martin, Motor cortex bilateral motor representation depends on subcortical and interhemispheric interactions. J. Neurosci. 29, 6196–6206 (2009).

20. C. L. R. Gonzalez, O. A. Gharbawie, P. T. Williams, J. A. Kleim, B. Kolb, I. Q. Whishaw, Evidence for bilateral control of skilled movements: Ipsilateral skilled forelimb reaching deficits and functional recovery in rats follow motor cortex and lateral frontal cortex lesions. Eur. J. Neurosci. 20, 3442–3452 (2004).

21. J.-Z. Guo, A. R. Graves, W. W. Guo, J. Zheng, A. Lee, J. Rodríguez-González, N. Li, J. J. Macklin, J. W. Phillips, B. D. Mensh, K. Branson, A. W. Hantman, Cortex commands the performance of skilled movement. eLife 4, e10774 (2015).

22. A. Miri, C. L. Warriner, J. S. Seely, G. F. Elsayed, J. P. Cunningham, M. M. Churchland, T. M. Jessell, Behaviorally selective engagement of short-latency effector pathways by motor cortex. Neuron 95, 683–696.e11 (2017).

23. A. Mathis, P. Mamidanna, K. M. Cury, T. Abe, V. N. Murthy, M. W. Mathis, M. Bethge, DeepLabCut: Markerless pose estimation of user-defined body parts with deep learning. Nat. Neurosci. 21, 1281–1289 (2018).

24. M. M. Churchland, A. Afshar, K. V. Shenoy, A central source of movement variability. Neuron 52, 1085–1096 (2006).

25. M. Hasegawa, K. Majima, T. Itokazu, T. Maki, U.-R. Albrecht, N. Castner, M. Izumo, K. Sohya, T. K. Sato, Y. Kamitani, T. R. Sato, Selective suppression of local circuits during movement preparation in the mouse motor cortex. Cell Rep. 18, 2676–2686 (2017).

26. D. B. Willingham, A neuropsychological theory of motor skill learning. Psychol. Rev. 105, 558–584 (1998).

27. R. E. Passingham, Attention to action. Philos. Trans. R. Soc. Lond. B Biol. Sci. 351, 1473–1479 (1996).

28. R. F. Baumeister, Choking under pressure: Self-consciousness and paradoxical effects of incentives on skillful performance. J. Pers. Soc. Psychol. 46, 610–620 (1984).

29. T. Krings, R. Töpper, H. Foltys, S. Erberich, R. Sparing, K. Willmes, A. Thron, Cortical activation patterns during complex motor tasks in piano players and control subjects. A functional magnetic resonance imaging study. Neurosci. Lett. 278, 189–193 (2000).

30. L. Ma, B. Wang, S. Narayana, E. Hazeltine, X. Chen, D. A. Robin, P. T. Fox, J. Xiong, Changes in regional activity are accompanied with changes in inter-regional connectivity during 4 weeks motor learning. Brain Res. 1318, 64–76 (2010).

31. D. S. Hoffman, P. L. Strick, Effects of a primary motor cortex lesion on step-tracking movements of the wrist. *J. Neurophysiol.* **73**, 891–895 (1995).

32. C. E. Lang, M. H. Schieber, Reduced muscle selectivity during individuated finger movements in humans after damage to the motor cortex or corticospinal tract. *J. Neurophysiol.* **91**, 1722–1733 (2004).

33. R. Kawai, T. Markman, R. Poddar, R. Ko, A. L. Fantana, A. K. Dhawale, A. R. Kampff, B. P. Ölveczky, Motor cortex is required for learning but not for executing a motor skill. *Neuron* **86**, 800–812 (2015).

34. O. Hikosaka, K. Nakamura, K. Sakai, H. Nakahara, Central mechanisms of motor skill learning. *Curr. Opin. Neurobiol.* **12**, 217–222 (2002).

35. H. H. Yin, B. J. Knowlton, The role of the basal ganglia in habit formation. *Nat. Rev. Neurosci.* **7**, 464–476 (2006).

36. H. Makino, E. Hwang, N. Hedrick, T. Komiyama, Circuit mechanisms of sensorimotor learning. *Neuron* **92**, 705–721 (2016).

37. T. Kitamura, S. K. Ogawa, D. S. Roy, T. Okuyama, M. D. Morrissey, L. M. Smith, R. L. Redondo, S. Tonegawa, Engrams and circuits crucial for systems consolidation of a memory. *Science* **356**, 73–78 (2017).

38. A. Mitani, T. Komiyama, Real-time processing of two-photon calcium imaging data including lateral motion artifact correction. *Front. Neuroinform.* **12**, 98 (2018).

39. L. Theis, P. Berens, E. Froudarakis, J. Reimer, M. Román Rosón, T. Baden, T. Euler, A. S. Tolias, M. Bethge, Benchmarking spike rate inference in population calcium imaging. *Neuron* **90**, 471–482 (2016).

40. B.-H. Liu, A. D. Huberman, M. Scanziani, Cortico-fugal output from visual cortex promotes plasticity of innate motor behaviour. *Nature* **538**, 383–387 (2016).

Citation: E. J. Hwang, J. E. Dahlen, Y. Y. Hu, K. Aguilar, B. Yu, M. Mukundan, A. Mitani, T. Komiyama, Disengagement of motor cortex from movement control during long-term learning. *Sci. Adv.* **5**, eaay0001 (2019).