Systems/Circuits

Disengagement of Motor Cortex during Long-Term Learning Tracks the Performance Level of Learned Movements

Eun Jung Hwang,* Jeffrey E. Dahlen,* Madan Mukundan, and Takaki Komiyama

Neurobiology Section, Center for Neural Circuits and Behavior, Department of Neurosciences, and Halıcıoğlu Data Science Institute, University of California San Diego, La Jolla, California 92093

Not all movements require the motor cortex for execution. Intriguingly, dependence on motor cortex of a given movement is not fixed, but instead can dynamically change over the course of long-term learning. For instance, rodent forelimb movements that initially require motor cortex can become independent of the motor cortex after an extended period of training. However, it remains unclear whether long-term neural changes rendering the motor cortex dispensable are a simple function of the training length. To address this issue, we trained mice (both male and female) to perform two distinct forelimb movements, forward versus downward reaches with a joystick, concomitantly over several weeks, and then compared the involvement of the motor cortex between the two movements. Most mice achieved different levels of motor performance between the two movements after long-term training. Of the two movements, the one that achieved higher trial-to-trial consistency (i.e., consistent-direction movement) was significantly less affected by inactivation of motor cortex than the other (i.e., variable-direction movement). Two-photon calcium imaging of motor cortical neurons revealed that the consistent-direction movement activates fewer neurons, producing weaker and less consistent population activity than the variable-direction movement. Together, the motor cortex was less engaged and less necessary for learned movements that achieved higher levels of consistency. Thus, the long-term reorganization of neural circuits that frees the motor cortex from the learned movement is not a mere function of training length. Rather, this reorganization tracks the level of motor performance that the animal achieves during training.

Key words: motor cortex inactivation; motor learning; mouse motor cortex; mouse reaching; movement consistency; multiple movement learning

---

Significance Statement

Long-term training of a movement reshapes motor circuits, disengaging motor cortex potentially for automatized execution of the learned movement. Acquiring new motor skills often involves learning of multiple movements (e.g., forehand and backhand strokes when learning tennis), but different movements do not always improve at the same time nor reach the same level of proficiency. Here we showed that the involvement of motor cortex after long-term training differs between similar yet distinct movements that reached different levels of expertise. Motor cortex was less engaged and less necessary for the more proficient movement. Thus, disengagement of motor cortex is not a simple function of training time, but instead tracks the level of expertise of a learned movement.

Introduction

The primary motor cortex (M1) is central to learning and execution of various skilled movements (Sanes and Donoghue, 2000; Graziano, 2006; Makino et al., 2016; Peters et al., 2017b; Papale and Hooks, 2018). However, not all movements require intact M1 for execution (Darling et al., 2011). For instance, species-specific innate motor behavior and/or gross movements, such as grooming, locomotion, and climbing, can spontaneously recover from lesion in M1 or its corticospinal output (Berridge and Whishaw, 1992; Muir and Whishaw, 1999; Lemon et al., 2012). Distinct movement modules, even if they are sequenced together to serve one behavioral goal, may differ in their M1 dependence. In a reach-to-grasp task, for example, forelimb reaching movements may be restored after M1 lesion, whereas grasping remains severely impaired (Lemon et al., 2012; Lemke et al., 2019). Confounding the matter, M1 dependence is not fixed

even for a given movement, but instead it can change across different contexts and learning stages. Inactivation of mouse M1 impairs whole-body turning when turns are induced by sensory stimulus, but not when mice make turns spontaneously (Heindorf et al., 2018). Mouse forelimb reaching movements with a joystick require M1 early in training, but not once the movements have been well learned and stereotyped (Hwang et al., 2019a).

In parallel to the longitudinal change of M1 dependence, the neural activity pattern in M1 dynamically changes during long-term learning. Most intriguingly, in layer 2/3 of the mouse M1, trial-to-trial consistency of population activity accompanying the learned movement gradually declines during the extended period of training after the motor skill is stereotyped, suggesting decoupling between M1 activity and the learned movement (Hwang et al., 2019a). In layer 5, the level of movement-related neural activity gradually decreases over the course of learning, selectively for corticostriatal neurons (Kupferschmidt et al., 2017). The dynamic nature of M1 engagement also agrees with longitudinal functional imaging of the human brain, which shows that the level of M1 activity associated with learned movements changes over the course of long-term motor learning (Ungerleider et al., 2002; Ma et al., 2010). Furthermore, perturbation of M1 activity in humans indicates that M1 is important for maintaining a learned motor skill during the early period of consolidation after learning, but not later, implying a transfer of motor memory from M1 to elsewhere with the passage of time during the consolidation period (Muellbacher et al., 2002).

A number of studies have found such changes in the M1 dependence of a learned movement during long-term learning, and parallel reorganization of motor circuits that may underlie that change. However, it remains elusive what are the factors that regulate M1 dependence. It has been proposed that M1 might provide a teaching signal, guiding plasticity in the downstream motor circuits, such as the striatum, so that the downstream areas can gradually take over control of the learned movements independent of M1 (Kawai et al., 2015; Wolff et al., 2019). Such a transfer of movement control from M1 to elsewhere might take place to allow highly reliable, autonomous execution of stereotyped movement patterns (Dhawale et al., 2019; Merel et al., 2019). If so, offloading motor control from M1 would be beneficial only after animals have learned to perform task-related movements with a high reliability. Thus, a potentially important factor that may regulate M1 disengagement might be the level of motor performance, such as accuracy, efficiency, and consistency of learned movements. In other words, the influence of M1 on the downstream motor circuits might be modified as a function of the achieved motor performance. This hypothesis makes a testable prediction that similar but distinct movements that are trained in parallel may exhibit different degrees of M1 dependence even after similar lengths of training if they attain different levels of motor performance. The current study tests this prediction and presents experimental results supporting the hypothesis that M1 involvement after long-term training depends on how well the animal performs the learned movement.

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 Institute of Health. Mice (6 weeks or older, male and female; calcium imaging: cross between *Gad2-IRES-Cre* [JAX 010802] and *Rosa26-CAG-LSL-tdTomato* [JAD007914]; optogenetic inactivation: *PV*-Cre [JAX 008069], or cross between *PV-Cre* and Ai32 [JAX 024109]) were housed in a room with a reversed light cycle (12 h-12 h). Experiments were performed during the dark period.

*Behavioral apparatus.* The behavioral apparatus was housed in a soundproof box (40 × 40 × 40 cm), and training sessions were performed in the dark. The components of the task included a joystick (M11L061P; CHProducts), a 17 inch computer monitor (for visual stimulus presentation, placed ~15 cm from the right eye of the mouse; E1713Sb; Dell), and a water port with photodiodes to sense licking. The joystick handle was custom-machined and fitted with a 1/16-inch thick brass rod that mice manipulated with their left forepaw. An electromagnet (EM050-3-222; APW) can mechanically immobilize the joystick at the origin. The joystick could be moved from the origin in any direction within the 90 degree range between forward (180°) and downward (270°) and had a dynamic range of 50 mm in each direction (see Fig. 1A, B). The 2D position of the joystick was continuously recorded at 1 kHz using a data acquisition card (USB6008; National Instruments) and custom MATLAB software. An open source real-time software package BControl (http://brodywiki.princeton.edu/bcontrol/) was used to control and record the task-sequence execution, stimulus selection, auditory cue presentation, and reward dispensing. The presentation of visual stimuli (100% contrast, full-field, square wave drifting gratings 0.04 cycles/degree and 3 cycles/s) was implemented using an open source MATLAB Psychtoolbox (http://psychtoolbox.org/).

Two-direction joystick task. In the two-direction joystick task, one of two orthogonal visual stimuli (forward or downward drifting gratings) was presented on the computer monitor. Two seconds after visual stimulus onset, an auditory cue (6 kHz pure tone) marked the answer period (up to 10 s) during which moving the joystick into the target zone (hereafter referred to simply as "target entry") led to a water reward. The visual stimulus and auditory cue remained on throughout the answer period. The target zone associated with the forward drifting gratings is the area spanning the polar angular space between 180° and 210° and the radial distance beyond a set threshold (6-10 mm; see Fig. 1B). The target zone for the downward gratings is the area spanning the angular space between 240° and 270° and the radial distance beyond the same set threshold. Errors (i.e., entering the incorrect target and movements before the "go" cue) triggered a white noise and led to an immediate trial termination. After reward, trial termination because of errors, or no response, and after the joystick returned to the origin, an intertrial interval (ITI, 4 s) began. During the ITI, the joystick was immovable, fixed at the origin by the electromagnet. At the end of the ITI (simultaneously with the beginning of the visual stimulus onset in the next trial), the electromagnet was disengaged and the joystick became movable.

Mice were trained under head fixation in the behavioral setup ~1 h per day over a period of 55 ± 3 d (mean ± SE across 22 mice). The task was shaped to reach the final version described above through three training blocks. In the first training block (10 ± 1 d), mice received a water reward as long as they moved the joystick to the correct target within a 30 s answer period (even if they hit the incorrect target first). As mice become proficient at moving the joystick in both directions, we increased the target distance from 6 mm up to 10 mm, and we decreased the response time to 10 s until they reached the targets during the answer period in >80% of trials.

In Block 2 (14 ± 2 d), mice were trained to refrain from moving the joystick before the go cue by immediately terminating a trial with a white noise error sound if they initiated movements before the go cue. Mice were rewarded if they successfully withheld their movement until the go cue and reached the correct target during the answer period. Entering the incorrect target before the correct target was not penalized. Block 2 training continued until mice achieved withholding performance >55%. Withholding performance continued to improve during the subsequent Block 3 training, reaching 73 ± 2.2% (mean ± SE across 22 mice) withholding during the last 11 training sessions.

In Block 3 (32 ± 4 d), trials were immediately terminated if mice moved before the go cue (as in Block 2) or reached the incorrect target. Training continued in Block 3 until correct choice performance reached 70%. The fraction of correct choice was computed for all trials that reached a target regardless of whether or not trials were successfully
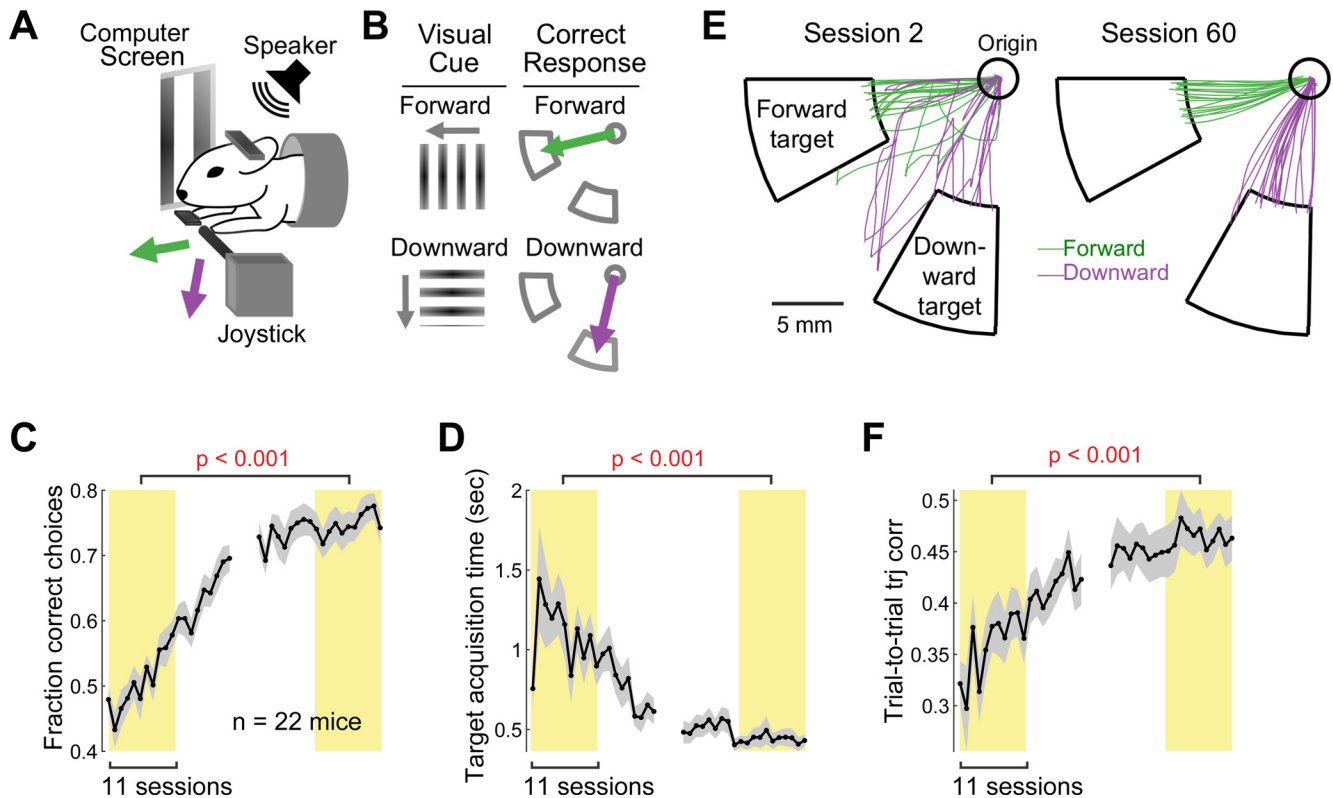
Figure 1. Long-term motor learning in the two-direction joystick task. ***A***, Behavioral task setup. A head-fixed mouse grabs the joystick with its left paw and makes forelimb reaches in the forward and downward directions. ***B***, Two-direction joystick task rule. Each trial begins with either forward or downward drifting gratings on the monitor, and the mouse must press the joystick in the same direction as the visual stimulus after an auditory tone to receive a reward. ***C***, Learning curve for the fraction of correct choices. Data are mean ± SE (*n* = 22 mice). The average fraction of correct choices in the first 11 training sessions was compared with the average fraction in the last 11 training sessions across 22 mice (paired bootstrap test; Table 1; see Materials and Methods). ***D***, Learning curve for the target acquisition time (i.e., duration between movement onset and target entry). The same illustration format and statistical test as in ***C***. ***E***, Example rewarded movements in early versus late training sessions (1 session/day). Green and purple traces represent forward and downward movements, respectively. Fifteen movements were randomly selected in each direction in each session. ***F***, Learning curve for movement consistency (i.e., the average trial-to-trial correlation coefficient between pairs of movement trajectories from rewarded trials in the same direction). The higher the correlation is, the more consistent the trajectories across trials are. The same illustration format and statistical test as in ***C***.

withheld. Once this performance criterion was met, the ITI length was increased up to 4 s (from 1 s). After achieving expertise during training (>55% of correct withholding and >70% of correct choice), mice began the optogenetic inactivation experiment or underwent a second surgery for the subsequent imaging experiment.

Inactivation experiment. An efficient way to suppress excitatory output from a brain region in selective trials is to photo-activate GABAergic neurons which potently inhibit local pyramidal neurons in the region (Li et al., 2019). A number of studies have used this indirect photoinhibition to successfully inactivate local cortical regions (Olsen et al., 2012; Z. V. Guo et al., 2014; Peters et al., 2014; J. Z. Guo et al., 2015; Hwang et al., 2019a). Hence, we opted to inactivate M1 by photoactivating parvalbumin (PV)-expressing GABAergic inhibitory neurons. Mice (either *PV*-Cre, *n* = 6; or *PV-Cre*::Ai32, *n* = 4) were implanted with a head fixation bar (i.e., headbar). *PV-Cre* mice were bilaterally injected with a virus carrying Channelrhodopsin-2 (ChR2; AAV2-1-EF1A-DIO-hChR2-eYFP, UPenn Vector Core) through a thinned skull window over the forelimb region of M1 (0.3 mm anterior to bregma, 1.5 mm lateral to the midline); ~100 nl of virus was injected at one location at each of two depths: 200 and 600 μm from the dura. The viral expression of ChR2 was typically confined within ~1 mm from the injection sites (see Fig. 3*A*,*B*). *PV-Cre*::Ai32 mice received bilateral glass windows over the forelimb region of M1. After surgery, mice followed the same training protocol outlined above before inactivation experiments. Once mice reached expertise, we inspected the previously thinned skull or glass windows for clarity, rethinning or replacing the glass window if necessary.

The inactivation experiment was performed across 10 daily sessions. Control and M1 inactivation sessions (see Fig. 3*C*) were alternated session-by-session, or in a block manner (e.g., 5 inactivation sessions

followed by 5 control sessions). In control sessions, the bifurcated LED fibers (Doric Lenses) were pointed above the headbar; whereas in M1 inactivation sessions, LED fibers were placed directly above M1 in both hemispheres (see Fig. 3*C*). Aside from LED fiber placement, all procedures were identical between control and M1 inactivation sessions.

In both control and inactivation sessions, we randomly selected ~15% of trials as light-on trials in which 465 nm blue light (i.e., light-on trials) was transmitted through the bifurcated LED fibers (7.5 mW per fiber). In light-on trials, light was continuously on from visual cue onset to the end of trial. A trial ended when the mouse made a premature movement before go cue (2 s after visual cue onset), when the mouse moved the joystick into any target within the answer period (10 s from go cue), or when the answer period was timed out. So, the light could stay on continuously for up to 12 s. Extracellular recording in our previous study confirmed that this approach with a continuous light effectively suppresses most of the neural activity in visual cortex where ChR2 was virally expressed in PV inhibitory neurons (Hwang et al., 2019a).

The two groups of mice that expressed ChR2 in PV neurons either virally or trans genically produced statistically indistinguishable results (see Fig. 5; Table 1), and we combined the two groups when presenting results in the main text.

Imaging experiment. After achieving expertise in the behavioral task, we temporarily suspended training and allowed mice unlimited access to water for at least 2 d before craniotomy and virus injections. The craniotomy spanned the forelimb region of M1 in the right hemisphere. Virus (AAV2-1-hSyn-GCaMP6f diluted in saline 1:8; UPenn Vector Core) was injected at 5 sites (~20 nl per site) in M1 at a depth of ~250 μm beneath the dura in layer 2/3. After injections, the craniotomy was covered with a glass window fixed in place with dental cement.

Table 1. Summary of statistical tests[a]

Figure	Test statistics and model variables	Test	95% CI
1C	$\Delta =$ Late − Early Fraction correct choices	Paired bootstrap	$0.20 < \Delta < 0.28$
1D	$\Delta =$ Late − Early Target acquisition time	Paired bootstrap	$-0.87 < \Delta < -0.46$
1F	$\Delta =$ Late − Early Trial-to-trial trajectory correlation	Paired bootstrap	$0.07 < \Delta < 0.16$
2D	$\Delta =$ Consistent − Variable Trial-to-trial trajectory correlation	Paired bootstrap	$0.11 < \Delta < 0.18$
2E	$\Delta =$ Consistent − Variable Target acquisition time	Paired bootstrap	$-0.23 < \Delta < -0.02$
2F	$\Delta =$ Consistent − Variable Fraction of correct choices	Paired bootstrap	$0.06 < \Delta < 0.20$
4B, panel 1	$x =$ Consistent or variable direction $y = \Delta$ no response	Mixed-effects model	$-0.01 <$ slope < 0.08
	$x =$ Inactivation or Control $y =$ No response, consistent direction	Mixed-effects model	$-0.03 <$ slope < 0.36
	$x =$ Inactivation or Control $y =$ No response, variable direction	Mixed-effects model	$-0.01 <$ slope < 0.41
4B, panel 2	$x =$ Consistent or variable direction $y = \Delta$ premature response	Mixed-effects model	$-0.10 <$ slope < 0.09
	$x =$ Inactivation or Control $y =$ Premature response, consistent direction	Mixed-effects model	$0.07 <$ slope < 0.46
	$x =$ Inactivation or Control $y =$ Premature response, variable direction	Mixed-effects model	$0.10 <$ slope < 0.41
4B, panel 3	$x =$ Consistent or variable direction $y = \Delta$ no target entry	Mixed-effects model	$-0.11 <$ slope < 0.10
	$x =$ Inactivation or Control $y =$ No target entry, consistent direction	Mixed-effects model	$0.14 <$ slope < 0.41
	$x =$ Inactivation or Control $y =$ No target entry, variable direction	Mixed-effects model	$0.14 <$ slope < 0.37
4B, panel 4	$x =$ Consistent or variable direction $y = \Delta$ incorrect target entry	Mixed-effects model	$0.10 <$ slope < 0.53
	$x =$ Inactivation or Control $y =$ Incorrect target entry, consistent direction	Mixed-effects model	$-0.06 <$ slope < 0.18
	$x =$ Inactivation or Control $y =$ Incorrect target entry, variable direction	Mixed-effects model	$0.20 <$ slope < 0.54
4C, panel 1	$x =$ Consistent or variable direction $y = \Delta$ reaction time	Mixed-effects model	$-0.65 <$ slope < 0.25
	$x =$ Inactivation or Control $y =$ Reaction time, consistent direction	Mixed-effects model	$-0.41 <$ slope < 1.33
	$x =$ Inactivation or Control $y =$ Reaction time, variable direction	Mixed-effects model	$-0.11 <$ slope < 0.53
4C, panel 2	$x =$ Consistent or variable direction $y = \Delta$ target acquisition time	Mixed-effects model	$0.01 <$ slope < 0.60
	$x =$ Inactivation or Control $y =$ Target acquisition time, consistent direction	Mixed-effects model	$-0.03 <$ slope < 0.43
	$x =$ Inactivation or Control $y =$ Target acquisition time, variable direction	Mixed-effects model	$0.20 <$ slope < 1.03
4C, panel 3	$x =$ Consistent or variable direction $y = \Delta$ trial-to-trial trajectory correlation	Mixed-effects model	$-0.12 <$ slope < 0.08
	$x =$ Inactivation or Control $y =$ Trial-to-trial trajectory correlation, consistent direction	Mixed-effects model	$-0.13 <$ slope < 0.04
	$x =$ Inactivation or Control $y =$ Trial-to-trial trajectory correlation, variable direction	Mixed-effects model	$-0.11 <$ slope < 0.03
4C, panel 4	$x =$ Consistent or variable direction $y = \Delta$ SD target acquisition time	Mixed-effects model	$0.06 <$ slope < 1.38
	$x =$ Inactivation or Control $y =$ SD target acquisition time, consistent direction	Mixed-effects model	$-0.10 <$ slope < 0.55
	$x =$ Inactivation or Control $y =$ SD target acquisition time, variable direction	Mixed-effects model	$0.23 <$ slope < 1.88
4C, panel 5	$x =$ Consistent or variable direction $y = \Delta$ path length	Mixed-effects model	$-0.35 <$ slope < 1.02
	$x =$ Inactivation or Control $y =$ Path length, consistent direction	Mixed-effects model	$-0.30 <$ slope < 0.16

(*Table continues.*)

Table 1. Continued

Figure	Test statistics and model variables	Test	95% CI
	x = Inactivation or Control	Mixed-effects model	$-0.41 <$ slope < 0.92
	y = Path length, variable direction		
4D, panel 1	Δ = Consistent − Variable	Paired bootstrap	$-0.10 < \Delta < 0.01$
	Δ no response		
	Δ = Inactivation − Control, consistent direction	Paired bootstrap	$0.01 < \Delta < 0.34$
	No response		
	Δ = Inactivation − Control, variable direction	Paired bootstrap	$0.20 < \Delta < 0.41$
	No response		
4D, panel 2	Δ = Consistent − Variable	Paired bootstrap	$-0.11 < \Delta < 0.06$
	Δ premature response		
	Δ = Inactivation − Control, consistent direction Premature response	Paired bootstrap	$0.02 < \Delta < 0.46$
	Δ = Inactivation − Control, variable direction Premature response	Paired bootstrap	$0.11 < \Delta < 0.44$
4D, panel 3	Δ = Consistent − Variable	Paired bootstrap	$-0.09 < \Delta < 0.15$
	Δ no target entry		
	Δ = Inactivation − Control, consistent direction	Paired bootstrap	$0.14 < \Delta < 0.46$
	No target entry		
	Δ = Inactivation − Control, variable direction	Paired bootstrap	$0.14 < \Delta < 0.36$
	No target entry		
4D, panel 4	Δ = Consistent − Variable	Paired bootstrap	$-0.54 < \Delta < -0.06$
	Δ incorrect target entry		
	Δ = Inactivation − Control, consistent direction Incorrect target entry	Paired bootstrap	$-0.08 < \Delta < 0.17$
	Δ = Inactivation − Control, variable direction	Paired bootstrap	$0.16 < \Delta < 0.49$
	Incorrect target entry		
4E, panel 1	Δ = Consistent − Variable	Paired bootstrap	$-0.27 < \Delta < 0.89$
	Δ reaction time		
	Δ = Inactivation − Control, consistent direction	Paired bootstrap	$-0.26 < \Delta < 1.39$
	Reaction time		
	Δ = Inactivation − Control, variable direction	Paired bootstrap	$-0.04 < \Delta < 0.91$
	Reaction time		
4E, panel 2	Δ = Consistent − Variable	Paired bootstrap	$-0.79 < \Delta < -0.11$
	Δ target acquisition time		
	Δ = Inactivation − Control, consistent direction	Paired bootstrap	$-0.05 < \Delta < 0.58$
	Target acquisition		
	Δ = Inactivation − Control, variable direction	Paired bootstrap	$0.32 < \Delta < 0.96$
	Target acquisition		
4E, panel 3	Δ = Consistent − Variable	Paired bootstrap	$-0.06 < \Delta < 0.19$
	Δ trial-to-trial trajectory correlation		
	Δ = Inactivation − Control, consistent direction	Paired bootstrap	$-0.13 < \Delta < 0.07$
	Trial-to-trial trajectory correlation		
	Δ = Inactivation − Control, variable direction	Paired bootstrap	$-0.18 < \Delta < 0.06$
	Trial-to-trial trajectory correlation		
4E, panel 4	Δ = Consistent − Variable	Paired bootstrap	$-1.46 < \Delta < -0.01$
	Δ SD target acquisition time		
	Δ = Inactivation − Control, consistent direction	Paired bootstrap	$-0.06 < \Delta < 0.85$
	SD target acquisition time		
	Δ = Inactivation − Control, variable direction	Paired bootstrap	$0.37 < \Delta < 1.61$
	SD target acquisition time		
4E, panel 5	Δ = Consistent − Variable	Paired bootstrap	$-0.83 < \Delta < 0.37$
	Δ path length		
	Δ = Inactivation − Control, consistent direction	Paired bootstrap	$-0.35 < \Delta < 0.28$
	Path length		
	Δ = Inactivation − Control, variable direction	Paired bootstrap	$-0.26 < \Delta < 0.69$
	Path length		
5A, panel 1	x = Consistent or variable direction	Mixed-effects model	$-0.02 <$ slope < 0.12
	$y = \Delta$ no response		
5A, panel 2	x = Consistent or variable direction	Mixed-effects model	$-0.13 <$ slope < 0.14
	$y = \Delta$ premature response		
5A, panel 3	x = Consistent or variable direction	Mixed-effects model	$-0.22 <$ slope < 0.12
	$y = \Delta$ no target entry		
5A, panel 4	x = Consistent or variable direction	Mixed-effects model	$0.28 <$ slope < 0.70
	$y = \Delta$ incorrect target entry		
5B, panel 1	x = Consistent or variable direction	Mixed-effects model	$-1.23 <$ slope < 0.34
	$y = \Delta$ reaction time		

(*Table continues.*)

Table 1. Continued

Figure	Test statistics and model variables	Test	95% CI
5B, panel 2	x = Consistent or variable direction y = Δ target acquisition time	Mixed-effects model	−0.20 < slope < 0.88
5B, panel 3	x = Consistent or variable direction y = Δ trial-to-trial trajectory correlation	Mixed-effects model	−0.10 < slope < 0.19
5B, panel 4	x = Consistent or variable direction y = Δ SD target acquisition time	Mixed-effects model	−0.21 < slope < 1.93
5B, panel 5	x = Consistent or variable direction y = Δ path length	Mixed-effects model	−1.00 < slope < 1.25
5C, panel 1	x = Consistent or variable direction y = Δ no response	Mixed-effects model	−0.02 < slope < 0.04
5C, panel 2	x = Consistent or variable direction y = Δ premature response	Mixed-effects model	−0.15 < slope < 0.12
5C, panel 3	x = Consistent or variable direction y = Δ no target entry	Mixed-effects model	−0.08 < slope < 0.17
5C, panel 4	x = Consistent or variable direction y = Δ incorrect target entry	Mixed-effects model	−0.27 < slope < 0.40
5D, panel 1	x = Consistent or variable direction y = Δ reaction time	Mixed-effects model	−0.75 < slope < 0.58
5D, panel 2	x = Consistent or variable direction y = Δ target acquisition time	Mixed-effects model	0.05 < slope < 0.42
5D, panel 3	x = Consistent or variable direction y = Δ trial-to-trial trajectory correlation	Mixed-effects model	−0.21 < slope < 0.06
5D, panel 4	x = Consistent or variable direction y = Δ SD target acquisition time	Mixed-effects model	0.11 < slope < 0.72
5D, panel 5	x = Consistent or variable direction y = Δ Path length	Mixed-effects model	0.20 < slope < 0.93
6A	x = Movement consistency y = Δ incorrect target entry	Mixed-effects model	−2.80 < slope < −1.37
6B	x = Movement consistency y = Δ target acquisition time	Mixed-effects model	−3.61 < slope < −0.12
6C	x = Movement consistency y = Δ SD target acquisition time	Mixed-effects model	−5.38 < slope < −1.26
7C, left	x = Consistent or variable direction y = Peak population average activity	Mixed-effects model	0.00 < slope < 0.01
7C, right	Δ = Consistent − Variable Peak population average activity	Paired bootstrap	0.00 < Δ < 0.03
7D, left	x = Consistent or variable direction y = Mean population average activity	Mixed-effects model	0.00 < slope < 0.01
7D, right	Δ = Consistent − Variable Mean population average activity	Paired bootstrap	0.00 < Δ < 0.01
7F, left	x = Consistent or variable direction y = Fraction of active neurons per trial	Mixed-effects model	0.00 < slope < 0.03
7F, right	Δ = Consistent − Variable Fraction of active neurons per trial	Paired bootstrap	0.01 < Δ < 0.04
7G, left	x = Consistent or variable direction y = Amplitude of active neurons per active trial	Mixed-effects model	−0.02 < slope < 0.01
7G, right	Δ = Consistent − Variable Amplitude of active neurons per active trial	Paired bootstrap	−0.03 < Δ < 0.01
7H, left	x = Consistent or variable direction y = Fraction of tuned neurons	Mixed-effects model	0.12 < slope < 0.22
7H, right	Δ = Consistent − Variable Fraction of tuned neurons	Paired bootstrap	0.10 < Δ < 0.22
7I, left	x = Consistent or variable direction y = Reliability of tuned neurons	Mixed-effects model	0.00 < slope < 0.04
7I, right	Δ = Consistent − Variable Reliability of tuned neurons	Paired bootstrap	0.01 < Δ < 0.04
8A, Left	x = Consistent or variable direction y = Fraction of active neurons per trial	Mixed-effects model	0.00 < slope < 0.04
8A, Right	x = Consistent or variable direction y = Fraction of active neurons per trial	Paired bootstrap	0.01 < Δ < 0.04
8A, Left	x = Consistent or variable direction y = Amplitude of activation neurons	Mixed-effects model	−0.02 < slope < 0.01
8A, Right	x = Consistent or variable direction y = Amplitude of activation neurons	Paired bootstrap	−0.01 < Δ < 0.01

(*Table continues.*)

Table 1. Continued

Figure	Test statistics and model variables	Test	95% CI
8A, Left	x = Consistent or variable direction y = Fraction of tuned neurons	Mixed-effects model	0.11 < slope < 0.22
8A, Right	x = Consistent or variable direction y = Fraction of tuned neurons	Paired bootstrap	0.04 < Δ < 0.2
8A, Left	x = Consistent or variable direction y = Reliability of tuned neurons	Mixed-effects model	0.00 < slope < 0.08
8A, Right	x = Consistent or variable direction y = Reliability of tuned neurons	Paired bootstrap	0.01 < Δ < 0.08
8B, left	x = Consistent or variable direction y = Fraction of active neurons per trial	Mixed-effects model	0.01 < slope < 0.03
8B, Right	x = Consistent or variable direction y = Fraction of active neurons per trial	Paired bootstrap	0.01 < Δ < 0.03
8B, Left	x = Consistent or variable direction y = Amplitude of activation neurons	Mixed-effects model	−0.01 < slope < 0.02
8B, Right	x = Consistent or variable direction y = Amplitude of activation neurons	Paired bootstrap	−0.01 < Δ < 0.02
8B, Left	x = Consistent or variable direction y = Fraction of tuned neurons	Mixed-effects model	0.11 < slope < 0.22
8B, Right	x = Consistent or variable direction y = Fraction of tuned neurons	Paired bootstrap	0.09 < Δ < 0.20
8B, Left	x = Consistent or variable direction y = Reliability of tuned neurons	Mixed-effects model	−0.00 < slope < 0.03
8B, Right	x = Consistent or variable direction y = Reliability of tuned neurons	Paired bootstrap	0.00 < Δ < 0.03
8C, left	x = Consistent or variable direction y = Fraction of active neurons per trial	Mixed-effects model	0.00 < slope < 0.03
8C, right	x = Consistent or variable direction y = Fraction of active neurons per trial	Paired bootstrap	0.01 < Δ < 0.04
8C, left	x = Consistent or variable direction y = Amplitude of activation neurons	Mixed-effects model	−0.02 < slope < 0.01
8C, right	x = Consistent or variable direction y = Amplitude of activation neurons	Paired bootstrap	−0.03 < Δ < 0.01
8C, left	x = Consistent or variable direction y = Fraction of tuned neurons	Mixed-effects model	0.08 < slope < 0.10
8C, right	x = Consistent or variable direction y = Fraction of tuned neurons	Paired bootstrap	0.05 < Δ < 0.18
8C, left	x = Consistent or variable direction y = Reliability of tuned neurons	Mixed-effects model	0.01 < slope < 0.01
8C, right	x = Consistent or variable direction y = Reliability of tuned neurons	Paired bootstrap	0.02 < Δ < 0.00
9B, left	x = Consistent or variable direction y = Trial-to-trial population activity correlation	Mixed-effects model	0.00 < slope < 0.02
9B, right	Δ = Consistent − Variable Trial-to-trial trajectory correlation	Paired bootstrap	0.00 < Δ < 0.01
10A	x = Movement consistency y = Peak population average activity	Mixed-effects model	−0.10 < slope < −0.04
10B	x = Movement consistency y = Fraction of active neuron per trial	Mixed-effects model	−0.28 < slope < −0.04
10C	x = Movement consistency y = Fraction of tuned neurons	Mixed-effects model	−1.34 < slope < −0.42
10D	x = Movement consistency y = Reliability of tuned neurons	Mixed-effects model	−0.23 < slope < −0.03
10E	x = Movement consistency y = Trial-to-trial population activity correlation	Mixed-effects model	−0.18 < slope < −0.03

[a]Test statistics are Δ in paired bootstrap tests and the slope in mixed-effects models. Plots in Figure 4 include three statistical tests, each corresponding to comparison of inactivation effects between consistent and variable direction, comparison of behavior between inactivation and control in the consistent direction, and comparison of behavior between inactivation and control in the variable direction.

Following surgery and recovery, mice resumed training until performance recovered to presurgery levels and GCaMP6f expression was sufficient to image (2-5 weeks after surgery). We imaged cortical activity in layer 2/3 at the depth of ~200 μm (see Fig. 7A) with excitation at 925 nm from a Ti-Sa laser (Mai Tai; Spectra-physics) using a two-photon resonant scanning microscope (B-scope; Thorlabs). Each imaging field was 512 × 512 pixels covering 472 × 508 μm and scanned at ~28.4 Hz. Each behavior-imaging session continued until mice became disengaged from the task, completed 170 rewarded trials, or the session reached 1 h,

whichever occurred first. Mice completed ~159 ± 30 (mean ± SE across 15 sessions) rewarded trials in each behavior-imaging session. We imaged 2 ± 1 fields per mouse (mean ± SE across 7 mice, range: 1-5), one field per session, across 8 mice. In our analyses that compared the two movement directions in relation to their asymmetric movement consistency, data from 1 mouse were excluded as this mouse showed nearly equal movement consistency between the two directions.

Movement analysis. Movement onset was defined as the first time at which the joystick velocity exceeded 20 mm/s continuously for 20 ms

and the joystick was displaced at least 1.1 mm from the origin. Target entry was defined as the first time at which the joystick entered a target zone. Movement offset was defined as the first time at which the joystick velocity fell below 20 mm/s continuously for 20 ms after target entry. A movement bout was bounded by movement onset and offset. The perimovement period was defined as temporal window from −0.5 to 2.25 s from movement onset, which included the whole movement bout for 99% of all successful trials. A small fraction of trials (4.0%) was excluded from our analyses, as movements in those trials progressed along the edges of the dynamic range of the joystick from the origin to the target, and we could not reliably estimate the animal's movement control in the direction perpendicular to the edge. Reaction time and target acquisition time were measured as the time from go cue to movement onset, and the time from movement onset to targe entry, respectively. Path length was measured as the cumulative displacement from movement onset to target entry.

Movement consistency: In each session, movement consistency in each direction was computed as the average trial-to-trial movement correlation across all pairs of movements that entered the correct target. Trial-to-trial movement correlation was measured using the Pearson's correlation coefficient between the two joystick traces (the concatenated *x* and *y* position time series) during the perimovement period.

Consistent versus variable direction: We computed movement consistency for each direction in the last 11 training sessions and designated the direction associated with a higher mean consistency as the consistent direction, and the other as the variable direction. The consistent direction identified this way exhibited a higher movement consistency in a majority of the 11 sessions in all but 1 mouse in the imaging experiment. In this exceptional mouse, movement consistency was nearly identical between the two directions and was higher for the consistent direction in 5 of 11 sessions. This mouse in the imaging experiment was excluded from all analyses in which we inferred differences related to asymmetric movement consistency between the two directions.

Neural activity analysis. Single-cell activity: Using custom MATLAB software, fluorescence images were aligned frame by frame to compensate for lateral motions. ROIs corresponding to individual neurons were manually drawn on the motion-corrected fluorescence images by circumscribing the cell bodies based on their GCaMP6f fluorescence intensity distinguishable from the background (mean: 160 neurons/session, range: 77-289). Pixels inside each ROI were considered as a single cell, whereas pixels extending radially outward from the cell boundary by 2-6 pixels were considered background. In the case the background included other cells' ROIs, those pixels were excluded. To estimate the activity of a single cell, 70% of the average pixel intensity in its background was subtracted from the average pixel intensity inside the cell (Hwang et al., 2017). The time series of the background-adjusted intensity was transformed to dF/F by dynamically estimating the baseline intensity (i.e., the eighth percentile of the intensity distribution in the 20 s window centered at each time point). dF/F was further transformed into spiking rate using the open source algorithm Suite2P (http://www.suite2p.org), in which the time constant of calcium sensor was set to 0.7 s.

Movement-related activation: A neuron was considered activated in a given trial during movement if its average activity during the perimovement period (−0.5 to 2.25 s from movement onset) exceeded a neuron-specific threshold. The threshold of each neuron was determined based on the distribution of its average activity during the baseline period (−3 to 0 s from cue onset). The baseline period covers the last part of the ITI during which mice are generally still. We present data analyses using mean + 2 SD of the baseline activity as the set threshold. However, results are robust across a wide range of thresholds (see Fig. 8A,B).

Activation reliability: For a given neuron, we computed the fraction of trials in which the neuron was activated during the perimovement period for each direction, respectively. A reliability of 1 indicates that the neuron is activated in every trial in which movement is made to the corresponding direction. A reliability close to zero indicates that the neuron is rarely activated.

Preferred direction of a neuron: For each neuron, the direction associated with a higher activation reliability was designated as its preferred direction. We also performed tuning analysis only using neurons that

showed reliability difference between the two directions of at least 0.01 and found similar trends to the results using all neurons (see Fig. 8C).

Trial-by-trial population activity correlation: The population activity correlation between two trials was measured as the correlation coefficient between the two concatenated perimovement activity time series of all neurons in the imaged population.

Statistical analysis. We collected data from multiple sessions for each animal in both inactivation and imaging experiments. As data (e.g., inactivation effects for consistent and variable direction) from the same session and from the same animal are not independent samples, we used linear mixed-effects models when assessing a statistically significant influence on a dependent variable *y* (e.g., inactivation effect) by an independent variable *x* (e.g., consistent vs variable direction) as follows. First, we fit our data with three mixed-effects models shown below, using the MATLAB function *fitlme*.

$$Model1 : y \sim 1 + x + (1|animal) + (x|animal)$$

$$Model2 : y \sim 1 + x + (1|animal) + (x|animal) + (x|animal : session)$$

$$Model3 : y \sim 1 + x + (1|animal) + (x|animal) + (1|animal : session)$$

In all three models, *y* is a linear function of *x*, but different random effect terms are added. In Model 1, random effects are added to account for the slope and the intercept that vary across animals. Models 2 and 3 include an extra random effect term that reflects dependence of data measured within each session. In Model 2, the extra random effect term accounts for the slope that varies across each animal's multiple sessions. In Model 3, the extra random effect term accounts for the intercept that varies across each animal's sessions. The two additional random effect terms in Models 2 and 3 could not be used simultaneously because of the limit in our sample size (i.e., the number of parameters become larger than the number of observations when including both terms). After fitting all three models, we selected the model that best explains our data as follows. We first compared Model 1 to Model 2 and to Model 3, using a likelihood ratio test as Model 1 is nested within the other two models. If neither Model 2 nor Model 3 is significantly better than Model 1 based on the likelihood ratio test, we selected Model 1 as the best model. Otherwise, we compared the log-likelihood between Models 2 and 3 and selected the one with a higher log-likelihood as the best model. If the fixed-effect slope in the best model is significantly different from zero ($p < 0.05$), we considered that there is a significant modulation of *y* by *x* (e.g., inactivation effect is significantly different between consistent and variable direction).

In addition to the mixed-effect model analysis, we performed a more conservative conventional significance test in which we computed the mean across multiple sessions per animal and performed nonparametric paired bootstrap test (Hwang et al., 2019b), to check the robustness of the mixed-effect model results.

The details of the test statistics and CIs of all statistical tests in the paper are described in Table 1.

Results

Long-term learning of the two-direction joystick task

To examine the relationship between motor performance and M1 involvement after long-term training, we trained mice to perform forelimb reaching movements in two distinct directions using a joystick for several weeks (55 ± 3 d, mean ± SE across 22 mice; 1 session/day). In this task, the mouse grabbed the joystick with their left paw and could move it from the origin in any direction between forward and downward directions (Fig. 1A). Each trial began with a visual stimulus, either forward or downward drifting gratings, on a monitor placed on the right side of the mouse. Two seconds after the visual cue onset, an auditory tone instructed the mouse to move the joystick (the "go" cue).
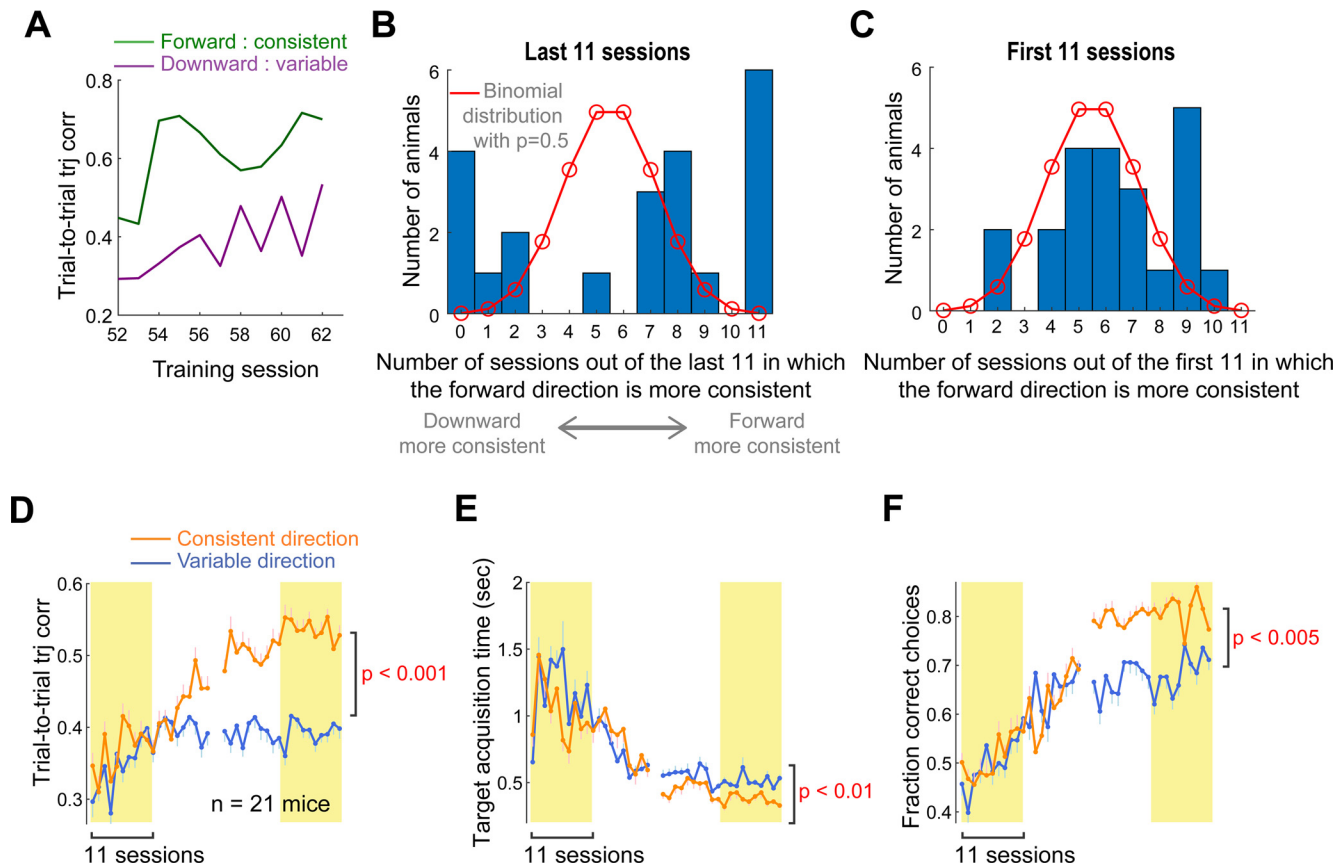
Figure 2. Asymmetric motor performance between the two reach directions. *A*, The trial-to-trial trajectory correlation in the last 11 training sessions of the example mouse shown in Figure 1*E*, separately for each direction. The forward movements are more consistent across trials than the downward movements, reliably across sessions in this mouse. *B*, The distribution of the number of sessions of the last 11 in which the trial-to-trial trajectory correlation is higher for movement in the forward than downward direction, across 22 mice. Red line indicates the expected, null distribution when the direction with a higher correlation is randomly chosen with a chance of 0.5 (i.e., binomial distribution). *C*, The distribution of the number of sessions of the first 11 in which trial-to-trial trajectory correlation is higher for movement in the forward than downward direction across 22 mice. The same illustration format as in *B*. *D*, The trial-to-trial trajectory correlation as a function of training session, respectively, for consistent and variable directions. The consistent direction of each mouse was defined as the direction associated with a higher trajectory correlation during the last 11 training sessions. Animals that exhibited the same consistent direction in at least 7 sessions of the last 11 were included. Data are mean ± SE ($n = 21$ mice). The mean correlation coefficient in the last 11 sessions was compared between the consistent and variable direction across 21 mice (paired bootstrap test; Table 1). *E*, The fraction of correct choices, respectively, for the consistent and variable directions defined solely based on the trial-to-trial trajectory correlation. The same illustration format and statistical test as in *D*. *F*, The target acquisition time. The same illustration format and statistical test as in *D*.

The mouse received a water reward if they moved the joystick beyond a set distance in the same direction as the visual stimulus ("target"; Fig. 1*B*). If the mouse moved the joystick before the go cue, moved the joystick into the incorrect target, or did not attain either target within the answer period (10 s), the trial was terminated without a reward.

Over the course of training, task performance improved gradually (Fig. 1*C*). The fraction of trials in which animals chose the correct target increased, indicating that they learned to associate the visual cues with the correct movement targets and learned to maneuver the joystick accurately to the correct targets. In parallel to the increased fraction of correct choices, the quality of movements that entered the correct targets (i.e., rewarded movements) also improved, similar to our previous observation in a single-direction joystick task (Hwang et al., 2019a). First, rewarded movements became more efficient as indicated by the decreased target acquisition time (i.e., time from movement onset to target entry) (Fig. 1*D*). Furthermore, rewarded movements became more consistent from trial to trial (Fig. 1*E*). We quantified the consistency of rewarded movements in each direction using the average correlation coefficient between the trajectories of those movements, and found that movement consistency significantly increased with training (Fig. 1*F*).

Different levels of consistency of the two learned movements
Although the overall task performance and movement quality improved after long-term training, the attained task performance and movement quality differed between the two directions. In most mice, movement consistency that was measured using between-trial trajectory correlation was higher for one direction than the other, reliably across sessions in the end of training (Fig. 2*A*). We examined movement consistency in the last 11 training sessions in each animal and classified which direction was more consistent (Fig. 2*A*). In 15 of 22 mice, movements in the forward direction were more consistent than the downward direction, and the opposite in 7 mice. To assess the robustness of the asymmetric movement consistency between the two directions, we examined the number of sessions of 11 in which the forward direction was more consistent for each animal. If the asymmetry reflects random variability across sessions, the number of sessions that show higher consistency for the forward direction would be close to 5 and 6 as expected by a binomial distribution with a chance of 0.5 for each direction. On the other hand, if the asymmetry reflects a true, persistent difference between the two directions in each animal, the number of sessions with higher consistency for the forward direction will be closer to 11 or 0. Compatible with the latter scenario, we found that the
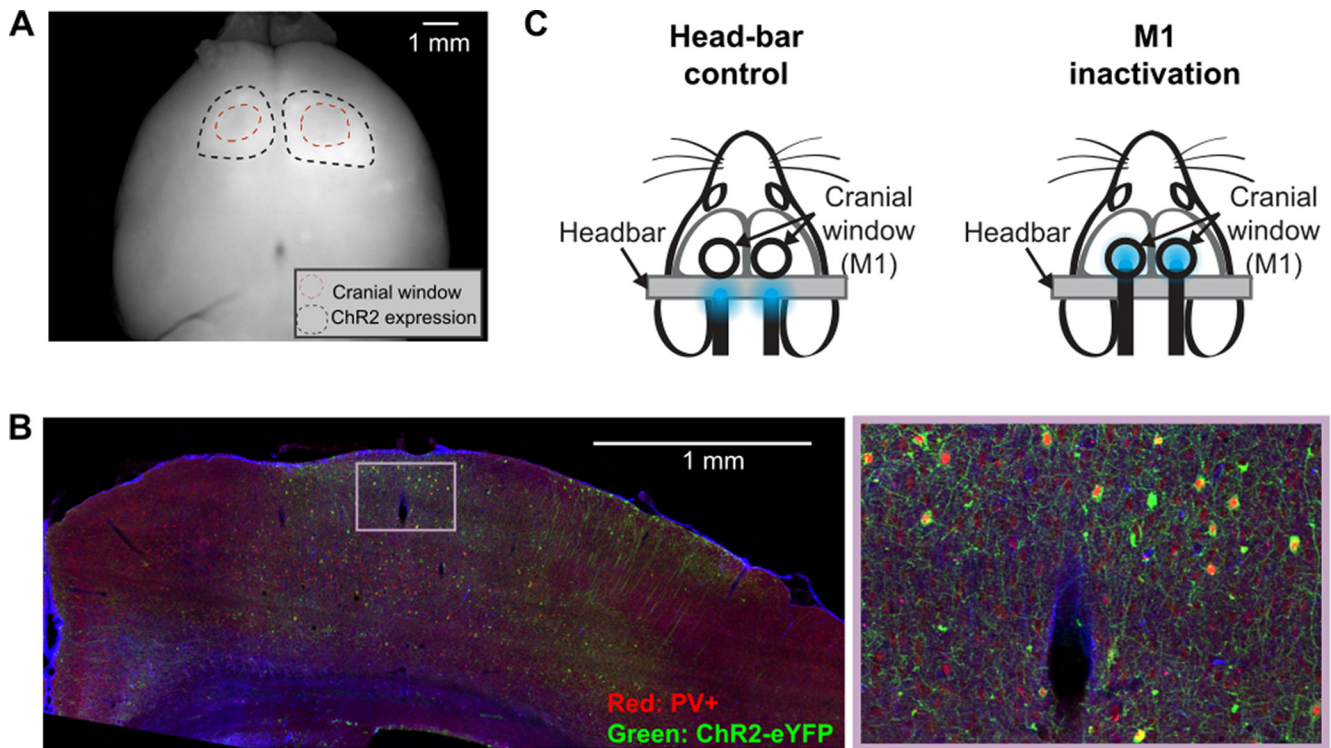
Figure 3. M1 inactivation experiment. *A*, A widefield fluorescence image of a whole brain extracted from a *PV-Cre* mouse injected with AAV2-1-EF1A-DIO-hChR2-eYFP in the motor cortex (M1). Black dotted lines indicate the boundary of ChR2-eYFP expression. Red dotted lines indicate the thinned skull areas over M1. *B*, A coronal slice near the forelimb region of M1 from a PV-Cre mouse injected with AAV2-1-EF1A-DIO-hChR2-eYFP (left) and a magnified view (right). PV$^+$ neurons are labeled with red fluorophore from a PV antibody staining. ChR2$^+$ neurons express eYFP. Most ChR2-expressing neurons are PV$^+$. *C*, Inactivation experiment setup. LED light is placed over the headbar bilaterally in control sessions, or over M1 in inactivation sessions. Each mouse performed 5 control and 5 inactivation sessions (1 session/day).

distribution is skewed to both ends, differently from the binomial distribution (Fig. 2B). Unlike the last training sessions, the equivalent analysis on the first 11 training sessions resulted in a distribution more similar to the binomial distribution (Fig. 2C), indicating that asymmetry in movement consistency had not been established in the early training days.

After defining the consistent and variable directions for each animal based on movement consistency in their last days of training, we examined the task performance and movement quality separately for the two directions. In this analysis, we removed 1 mouse which showed a weak asymmetry as a consistent direction was not robust across sessions (i.e., the movement consistency was higher for the forward direction in 5 of 11 sessions). However, results described below remain valid even when this mouse is included. By definition, movement consistency was significantly higher in the consistent than variable direction during the last 11 training sessions (Fig. 2D). However, it is noteworthy that the consistent direction defined from the last 11 training sessions had already started exhibiting significantly higher movement consistency after session 18, indicating that the asymmetry gradually developed and was maintained over the course of training. In addition, the target acquisition time was significantly shorter for the consistent than variable direction (Fig. 2E). Thus, movements in the consistent direction were not only more similar across trials, but also more efficient in achieving the goal than movements in the variable direction. We also compared choice behavior in response to the visual cues associated with each direction and found that the fraction of correct choice was significantly higher for the cue associated with the consistent direction than of the variable (Fig. 2F). The number of trials performed in

each direction over the course of training was not significantly different between the two directions (3419 ± 228 vs 3302 ± 267 for consistent vs variable direction; mean ± SE across 22 mice; paired bootstrap test, $p = 0.25$). Thus, the amount of training likely does not explain the observed behavioral differences between the two directions. At present, we do not know the exact origins of the behavioral differences. Nonetheless, in most mice, we found robust asymmetry in task performance and movement quality between the two directions after long-term training.

Variable movements are more dependent on M1
To examine the relationship between M1 dependence and motor performance, we conducted inactivation experiments in 10 expert mice. To inactivate M1, we activated PV-expressing inhibitory interneurons that express ChR2 (see Materials and Methods; Fig. 3A,B) by shining blue light into the cranial windows over the forelimb region of M1 bilaterally (Fig. 3C). Each mouse was also subjected to a control experiment in separate behavioral sessions in which light was applied to their headbar away from the cranial windows, but otherwise identical to the inactivation experiment (Fig. 3C). To remove any potential non-specific effects, such as distraction, because of the mice seeing the stimulation light, all inactivation effects reported hereafter were assessed by comparing the headbar control and M1 inactivation trials. The statistical significance of inactivation effect on any variable of interest was assessed using a linear mixed-effects model in which the variable of interest was modeled as a linear function of the treatment (i.e., control or inactivation) with random effects across animals, and across animals and sessions (see Materials and Methods). If the fixed-effect slope of the model is
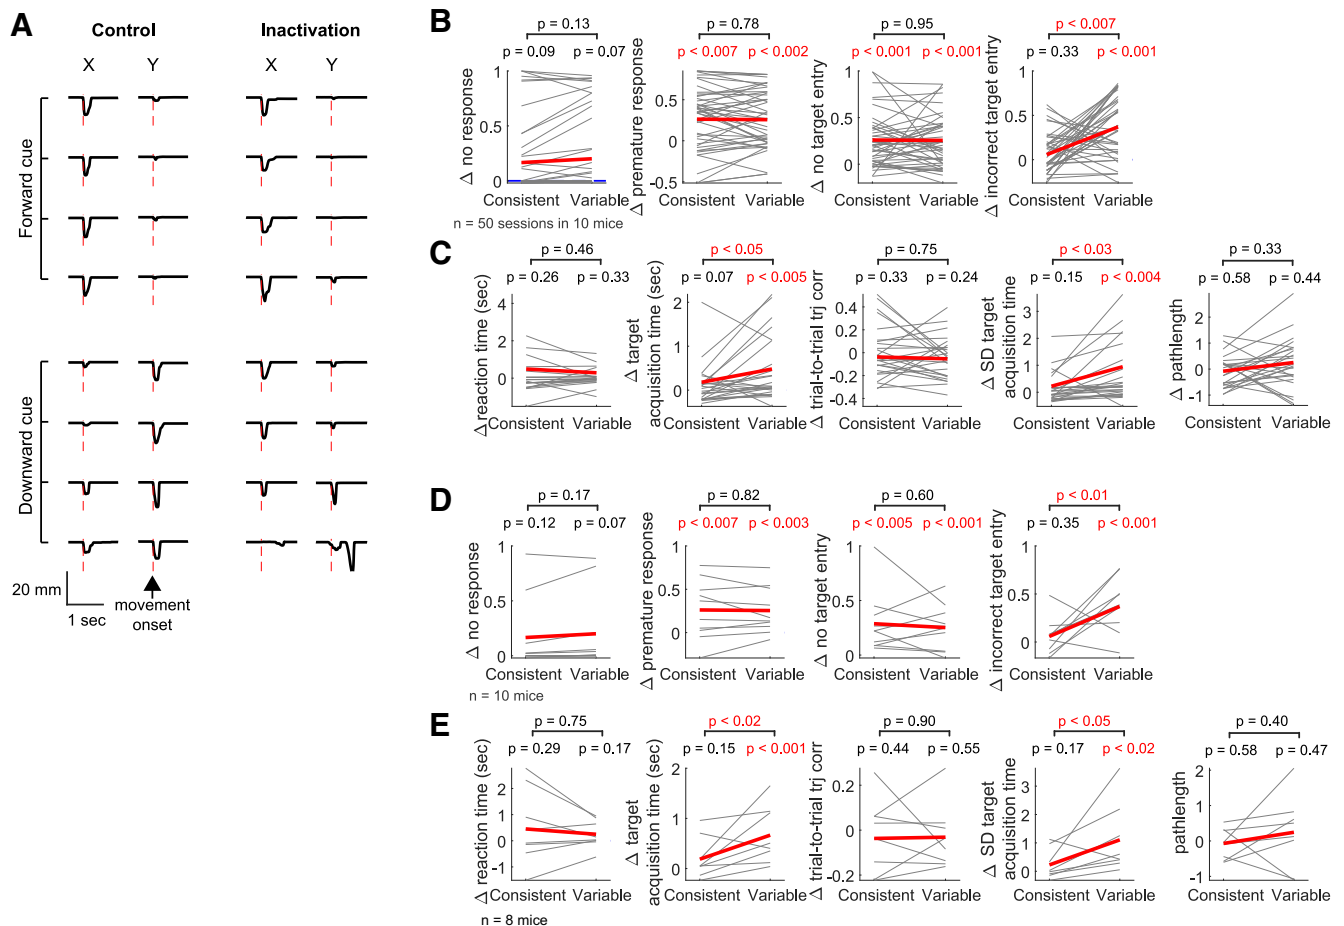
Figure 4. Asymmetric effects of M1 inactivation between the two reach directions. **A**, Movement trajectories in a control (left) versus inactivation (right) session. Negative displacements along the x axis correspond to movements in the forward direction, whereas negative displacements along the y axis correspond to the downward direction. Top (bottom) traces are movements in response to the forward (downward) drifting gratings. Red vertical lines indicate movement onset. There are incorrect choices in response to the downward drifting gratings in inactivation trials. In this mouse, movements in the forward direction were more consistent than in the downward direction during the last training sessions before the inactivation experiment. **B**, The effect of M1 inactivation (Δ = inactivation − control) on the fraction of errors. Four types of errors are examined: no response, premature response, no target entry, and incorrect target entry. The consistent and variable directions of each mouse are determined from their last 11 training sessions before the inactivation experiment. Thin lines indicate individual sessions. The direction dependence of inactivation effect was estimated from the slope of a linear mixed-effect model in which the direction is the independent variable and inactivation effect is the dependent variable (see Materials and Methods; Table 1). The red line and pink shade in each plot represent the 95% CI of the fixed-effect slope, and the statistical significance of the slope is also specified. A significant slope indicates that the inactivation effect shows a significant direction dependence. In the fraction of no response, 19 of 50 individual sessions are overlapped as their inactivation effects were zero (i.e., $\Delta = 0$) for both directions. **C**, The effect of M1 inactivation on the kinematics of movements that successfully entered the correct target. The same illustration format and statistical tests as in **B**. **D**, The same as in **B**, but per animal basis. The average effect was compared between the consistent and variable direction across 10 mice (paired bootstrap test). **E**, The same as in **C**, but per animal basis.

significantly different from zero (i.e., the variable of interest is significantly modulated by the treatment), we declared that the inactivation effect is significant. The mixed-effect model allows us to account for the variables measured from multiple experimental sessions in each mouse without collapsing all the sessions from each animal by averaging. We also compared inactivation effects (Δ = inactivation − control) between the consistent and variable directions, using a similar mixed-effects model in which the inactivation effect was modeled as a linear function of the movement direction with the random effects described above.

Figure 4A displays example movement trajectories in control versus inactivation condition. In the control condition, this mouse made the correct directional movements as instructed by the visual cue for both directions, and movements in the forward direction were more consistent across trials (Fig. 4A, left column). Inactivation of M1 did not generally alter consistent-direction (i.e., forward) movements. In contrast, compared with the control condition, response to the cue associated with the

variable direction was often erratic when M1 was inactivated (Fig. 4A, right column). To systematically examine inactivation effects on task performance, we quantified the fractions of errors in response to the cues instructing the consistent and variable directions. The consistent and variable directions of each mouse were defined in the same way as described in the previous section, using their last 11 training sessions before the inactivation experiment started. Four different types of errors were assessed in this task: (1) no response, (2) premature response before the go cue, (3) response but no entry to either target, and (4) entry to the incorrect target. Single trials can be assigned to more than one error type, such as error Types 2 and 3 or error Types 2 and 4. We found that M1 inactivation induces only marginal changes in the fraction of no response in both directions (Fig. 4B). In contrast, inactivation significantly increased the fraction of premature response in both directions, and the size of inactivation effect was not significantly different between the two directions (Fig. 4B). The increased fraction of premature response appears

to be consistent with the view that motor cortex is involved in suppressing movements (Ebbesen and Brecht, 2017). M1 inactivation also significantly increased the fraction of trials in which movement was initiated but failed to enter any target, and the effect size was not significantly different between the directions (Fig. 4B), suggesting that M1 inactivation impairs the animals' ability to complete learned movements in this task. Nevertheless, in a majority of inactivation trials (60 ± 16%; mean ± SE across 10 mice), movements reached the targets. Of those target entry trials, we examined the fraction of incorrect target entry and found that inactivation significantly increased incorrect target entry in response to the cue instructing the variable direction but not the consistent (Fig. 4B). In other words, animals under M1 inactivation erred more frequently by moving to the consistent-direction target in response to the cue instructing the variable direction but not vice versa. These results, based on the linear mixed-effects model that took into account all individual session data, were reproduced in the conventional per-animal basis analysis (Fig. 4D). Therefore, M1 inactivation induced deficits in withholding and completing learned movements in response to the cue for both directions, suggesting that M1 plays a role in suppressing movements and maintaining the initiated motor program. In addition, M1 inactivation increased the fraction of incorrect target entry selectively when mice were required to make movements in the variable direction, suggesting that the production of the variable-direction movement relies on M1 more strongly than the consistent-direction movement.

To further examine the M1 dependence of learned movements, we compared inactivation effects on kinematics of the movements that successfully entered the correct target between the two directions. We found that inactivation effects on the mean and variability of target acquisition time were selectively significant in the variable direction but not the consistent, and the effect sizes were significantly larger for the variable direction (Fig. 4C,E). That is, compared with the control condition, movements in the variable direction entered the target more slowly and the target acquisition time was more variable across trials when M1 was inactivated. On the other hand, the kinematics of movements in the consistent direction were not significantly affected by M1 inactivation. We found that these inactivation results are similar between the two groups of mice, virally versus transgenically expressing Ch2 in PV neurons (Fig. 5). Therefore, M1 appears to be more strongly involved in the kinematic control of movements in the variable direction than in the consistent.

Our analyses thus far compared inactivation effects in a binary categorical manner (i.e., consistent vs variable) and found significantly larger inactivation effects on task performance and movement kinematics in the variable direction than the consistent. This finding poses the question of whether M1 dependence of learned movements is modulated by movement consistency at a finer scale beyond the binary manner comparison. To test this idea, we used the raw measures of movement consistency for each mouse as the independent variable of the aforementioned linear mixed-effect model, instead of the binarized movement direction (see Materials and Methods). We found significantly negative slopes in the mixed-effect model for the fraction of incorrect target entry, the mean target acquisition time, and the variability of target acquisition time (Fig. 6). That is, inactivation effects on these variables were significantly modulated by the continuously varying movement consistency such that movements are progressively less affected by M1 inactivation as movement consistency increases.

Optical illumination heats brain tissue, which could suppress spiking activity in the region and thus alter behavior even in the absence of opsins (Owen et al., 2019). Our inactivation results might have been mediated by such nonspecific effects of light in M1. However, regardless of the exact mechanism underlying the perturbation of M1 activity, our interpretation still holds that the variable-direction movement is more strongly affected by M1 activity perturbation than the consistent-direction movement.

Relationship between learned movement consistency and M1 activation level

To investigate how M1 activity relates to movement consistency, we recorded the activity of neurons in M1 while expert mice performed the two-direction task (n = 7 mice). We virally expressed the genetically encoded calcium indicator GCaMP6f (AAV2-1-hSyn-GCaMP6f) in the forelimb regions of M1 and imaged the activity of individual neurons using two-photon microscopy (Fig. 7A).

We first compared the population average activity (i.e., activity averaged across all neurons in each imaging field) during rewarded movements that successfully entered the correct target between the consistent and variable directions. The consistent and variable directions of each mouse were identified based on their last 11 training sessions before the imaging experiment began, identically to the inactivation experiment. We found that the population average activity was stronger when mice made movements to their variable direction than to their consistent (Fig. 7B). Both the peak and mean amplitudes of the population average activity during the perimovement period were significantly larger for the variable direction than the consistent (Fig. 7C,D).

The stronger population average activity during movements in the variable direction may arise from multiple origins, namely, (1) the number of neurons activated during each trial is greater for the variable direction than for the consistent, and/or 2) the activity of each neuron activated during movement is stronger for the variable direction than for the consistent. To test the first possibility, we classified whether each neuron was activated or not during movement in each trial (Fig. 7E). A neuron was considered activated in a given trial if its perimovement period activity exceeded a threshold determined from the distribution of its baseline period activity (mean ± 2 SDs; see Materials and Methods). We found that a significantly larger fraction of neurons was activated per trial when mice made movements in the variable than consistent direction (Fig. 7F). To examine whether the second possibility also contributes to the observed difference in the strength of population activity, we compared the amplitude of perimovement activity of activated neurons in each trial between the two directions. Incompatible with the second possibility, we found that the amplitude of the activated neurons was not different between the two directions (Fig. 7G). Thus, the stronger population average activity for the variable direction is mainly driven by a larger fraction of M1 neurons being activated when mice make movements in the variable direction than the consistent.

To further investigate the different levels of M1 activation between the two directions, we characterized the activation reliability and preferred direction for each neuron. The activation reliability is defined as the fraction of trials in which each neuron is activated, and the preferred direction is the movement direction associated with a higher reliability (see Materials and Methods). We found that a significantly larger fraction of neurons prefers the variable direction (Fig. 7H). That is, a larger

PV-Ch2 Virally expressed mice

PV-Ch2 Transgenic mice
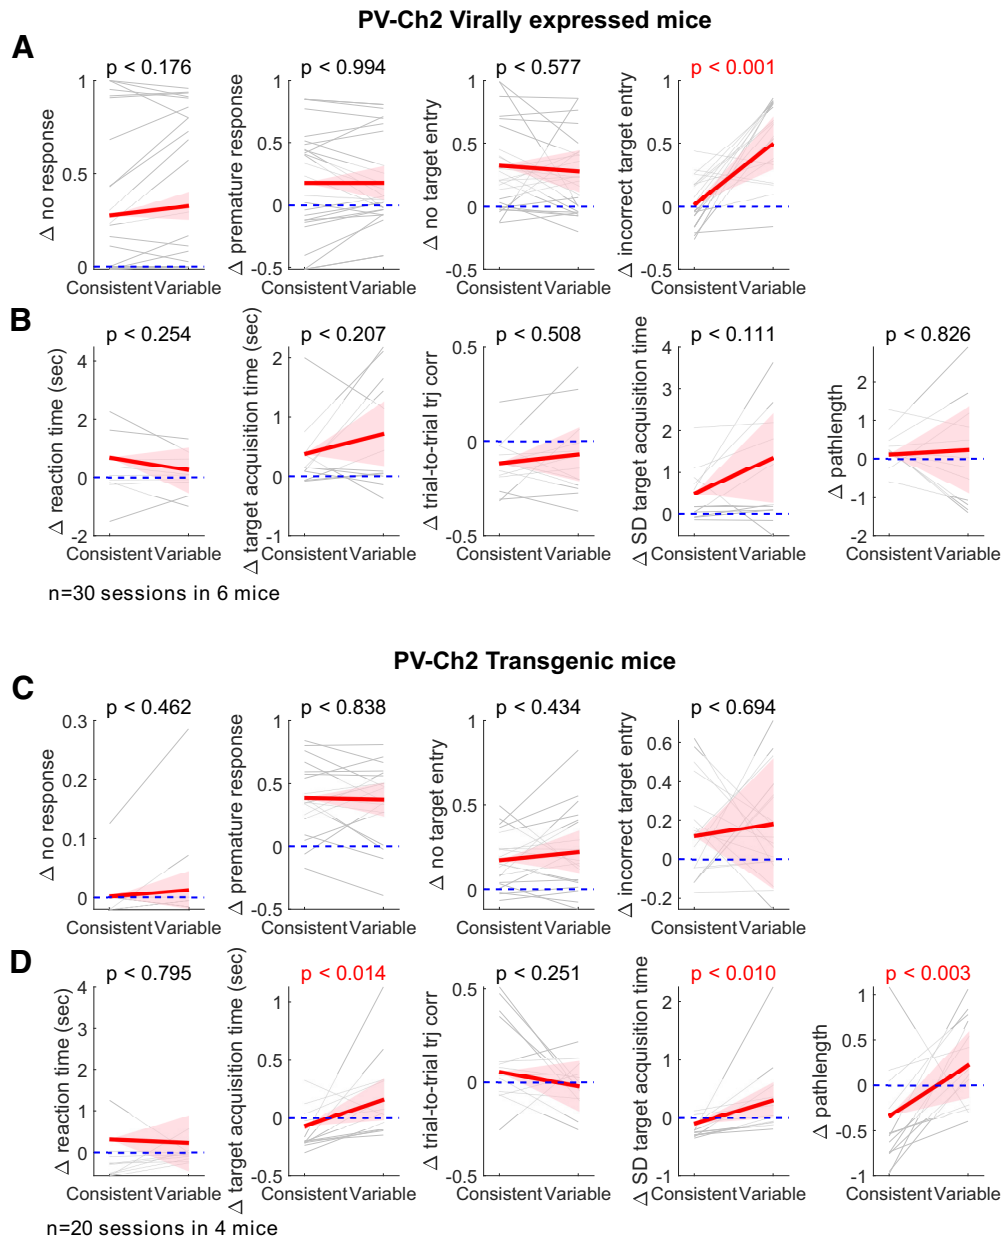
Figure 5. M1 inactivation effects, separately analyzed for two experimental groups. *A*, *B*, The same analyses as Figure 4*B*, *C*, but only with 6 mice in which ChR2 were expressed in PV neurons virally. *C*, *D*, The same analyses as Figure 4*B*, *C*, but only with 4 transgenic mice (*PV-Cre*::Ai32).

fraction of neurons is activated more frequently for the variable direction. Furthermore, the activation reliability across trials in the preferred direction was higher for neurons which preferred the variable direction than for those of the consistent (Fig. 7*I*). We found that the neuronal activity difference between the consistent and variable direction movements is robust across different activation thresholds and tuning strengths (Fig. 8). These results together indicate that a larger fraction of M1 neurons is preferentially and more reliably activated during movements in the variable direction, resulting in a larger overall fraction of activated M1 neurons and thus larger ensemble average per trial for the variable direction.

Relationship between learned movement consistency and M1 activity consistency

The higher activation reliability of neurons preferring the variable direction is unexpected given that movement trajectories to the variable direction are by definition less consistent across trials. So we asked whether this counterintuitive activity pattern extends to the reliability at the population activity level. Figure 9*A* displays example perimovement population activity from individual trials in the consistent versus variable directions in 1 animal, showing more consistent population activity in the variable direction than the consistent, similar to the single-neuron reliability. To systemically compare the reliability of population activity across trials, we quantified the trial-to-trial consistency of population activity using a correlation coefficient separately for each direction (see Materials and Methods). We found that the population activity is indeed less consistent during movements in the consistent direction than in the variable (Fig. 9*B*).

These differential activity patterns between the consistent and variable directions suggest that the engagement and consistency of the M1 population may vary at a finer scale according to the continuously varying movement consistency. To test this idea,
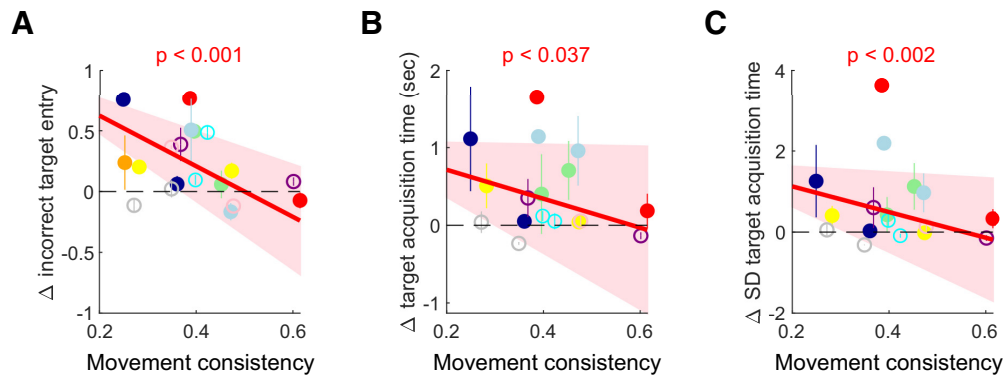
Figure 6. M1 inactivation effects modulated by the consistency of learned movements. **A**, The effect of M1 inactivation on the fraction of incorrect target entry as a function of the movement consistency (i.e., the trial-to-trial trajectory correlation) of each animal that was computed from their last 11 training sessions for each direction separately. Different colors for different mice. For illustration purposes, inactivation effects for each animal were shown as mean ± SE across their inactivation sessions. Solid circles represent *PV-Cre* mice with virally expressed ChR2. Empty circles represent *PV-Cre*::Ai32 mice. The dependence of inactivation effect on the movement consistency was estimated using a linear mixed-effect model, similar to Figure 4B, but movement consistency that continuously varied across animals was used as the independent variable. **B**, The effect of M1 inactivation on the mean target acquisition time as a function of movement consistency. The same format and statistical test as in **A**. **C**, The effect of M1 inactivation on the variability of target acquisition time as a function of movement consistency. The same format and statistical test as in **A**. Inactivation effects on the fraction of incorrect target entry, mean target acquisition time, and variability of target acquisition time showed significantly negative slopes (Table 1), indicating that effects decrease proportionally with movement consistency.

we examined various features of population activity as a function of the continuous varying movement consistency across mice. These activity features included the peak amplitude of population average activity, the fraction of activated neurons per trial, the fraction of directionally tuned neurons, the reliability of directionally tuned neurons, and the trial-to-trial population activity correlation (Fig. 10A–E). In support of our prediction, we found significantly negative slopes in all five variables, indicating that the engagement of the M1 population and their trial-to-trial consistency gradually decrease as the consistency of learned movements increases. Therefore, in addition to movement dependence on M1, the engagement of the M1 population also changes during long-term training in accordance with the increasing consistency of the learned movement. The counterintuitive finding that the M1 population is more consistent during more variable movements further highlights the notion that M1 population disengages gradually as a learned movement becomes highly consistent.

Discussion

Learning new motor skills often involves learning of multiple movement components (e.g., forehand and backhand strokes when learning tennis), but different movements do not always improve at the same time nor reach the same level of proficiency. Analogously, we trained mice to perform forward and downward forelimb reaching movements concomitantly over several weeks, but the two-directional movements attained different levels of motor performance. Intriguingly, the asymmetric levels of motor consistency between the two-directional movements revealed that motor consistency is a critical factor influencing long-term changes in neural circuits that control the learned movement. Within each animal, the more consistent movements were less affected by M1 inactivation than the variable direction and accompanied by weaker and less consistent M1 population activity. Across animals, the M1 population activity and inactivation effects varied with each animal's motor consistency; the more consistent the learned movement is, the less engaged and less necessary M1 is. Thus, M1 dependence and engagement are not a mere function of training time. Instead, movements that acquire higher levels of expertise become more disconnected from M1 control.

Relationship between movement consistency and M1 activity consistency

The higher trial-to-trial variability of M1 population activity during more consistent movements appears counterintuitive. However, the seemingly counterintuitive result may be explained by our previous finding from a single-direction forelimb reach task (Hwang et al., 2019a). In that study, we found that the learned movement becomes highly consistent and independent of M1 after long-term learning, but the corresponding population activity in M1 becomes less consistent from trial to trial, compared with the middle stage of learning (Fig. 10F). This long-term change may reflect rewiring among neurons in M1, influencing their correlative firing patterns (Meamardoost et al., 2020). Considering the previous finding, the two different movements in the current study might be at different learning stages. The animals were learning the variable-direction movement more slowly, such that our experiments might be timed at their mid-stage learning (Fig. 10G). In contrast, the consistent-direction movement was learned more quickly; thus, the experiments were done closer to late-stage learning (Fig. 10F). Accordingly, the consistent movement would be more strongly decoupled from the M1 activity, which allows M1 activity to be more variable across trials than the variable movement.

Learning-related activity changes in different layers of M1

Although our inactivation experiments silenced M1 activity presumably throughout the cortical layers, our imaging experiments examined neuronal activity only in layer 2/3. Layer 2/3 and layer 5 show different patterns of learning-related changes (Masamizu et al., 2014; Peters et al., 2017a), and inhibiting layer 5 activity can impair expert-level motor performance in certain tasks (Biane et al., 2019; Sauerbrei et al., 2020). However, the involvement of layer 5 at the expert stage might be task-dependent as learning-related disengagement has also been found in layer 5 neurons. The level of movement-related activity of M1 layer 5 neurons that project to the dorsolateral striatum (DLS) gradually decreases over the course of motor skill learning (Kupferschmidt et al., 2017). Furthermore, selective blockade of synaptic transmission from the DLS-projecting M1 neurons during early, but not late, training impairs motor learning, indicating an early
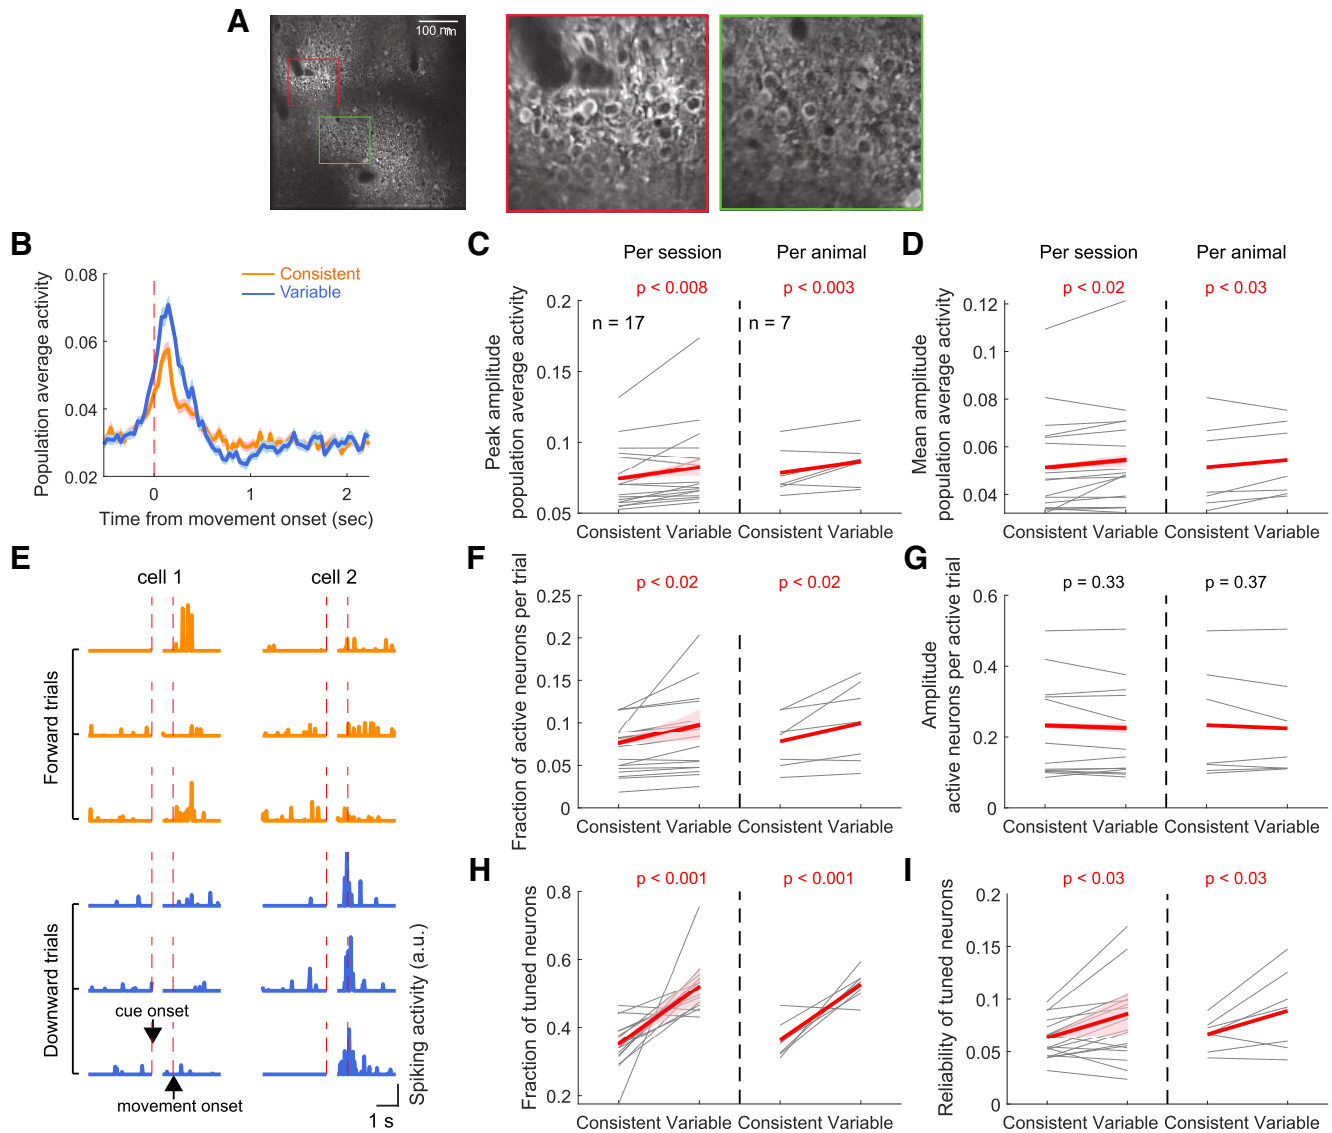
Figure 7. M1 neural activity modulated by the consistency of learned movements. **A**, The session-average fluorescence image of an example whole field of view in M1. Right two images represent magnified views showing individual neurons expressing GCaMP6f. **B**, The population average activity during movements in the consistent versus variable directions in an example M1 field. Data are mean ± SE across trials. The consistent and variable directions of each mouse were defined from their last training sessions before the imaging experiment. **C**, The peak amplitude of perimovement, population average activity during movements in the consistent versus variable directions. Left, The movement direction dependence of the peak amplitude was assessed using a mixed-effect model similar to Figure 4B. Red line and pink shade in each plot represent the 95% CI of the slope, and the statistical significance of the slope is specified (Table 1). Significant slopes indicate a significant dependence of the activity amplitude on the movement direction. Right, The same as at left, but per animal basis analysis. The average across sessions was taken for each animal. The amplitude was compared between the consistent and variable directions across 7 mice (paired bootstrap test). The same illustration format and statistical tests were applied to **C–D** and **F–I**. **D**, The mean amplitude of perimovement, population average activity. **E**, The activity of two example neurons in three forward and three downward movement trials. In each plot, the left side of the activity is aligned to cue onset, whereas the right side is aligned to movement onset. Cell 1 is tuned to the forward direction (consistent direction) as it is more reliably activated during the perimovement period in forward movement trials. Cell 2 is tuned to the downward direction. **F**, The fraction of neurons activated per trial when movement is made to the consistent versus variable direction. **G**, The perimovement activity amplitude of neurons that are activated. **H**, The fraction of neurons tuned to the consistent versus variable directions. **I**, The reliability of activation during movement to the preferred direction in consistent direction versus variable direction tuned neurons.

transient role of the corticostriatal projection neurons in motor skill learning (Wolff et al., 2019). Thus, in some tasks, corticostriatal neurons in M1 layer 5 may also become decoupled from the learned movement once the learned motor skill is consolidated.

The origins of different motor performance between the two-directional movements

The two forelimb movements in our task attained different levels of motor performance despite the almost equal durations and

trial numbers of training. The direction with better performance was idiosyncratic such that some animals were better in the forward direction versus other animals that preferred the downward direction. A multitude of factors may contribute to the unequal performance. For instance, dynamics of the joystick device (i.e., relationship between the motion and torque) are not perfectly symmetric between the two directions, possibly rendering one direction easier to control than the other. The biomechanics involved in moving the forelimb (e.g., the used muscle groups and their coordination) are also intrinsically different between
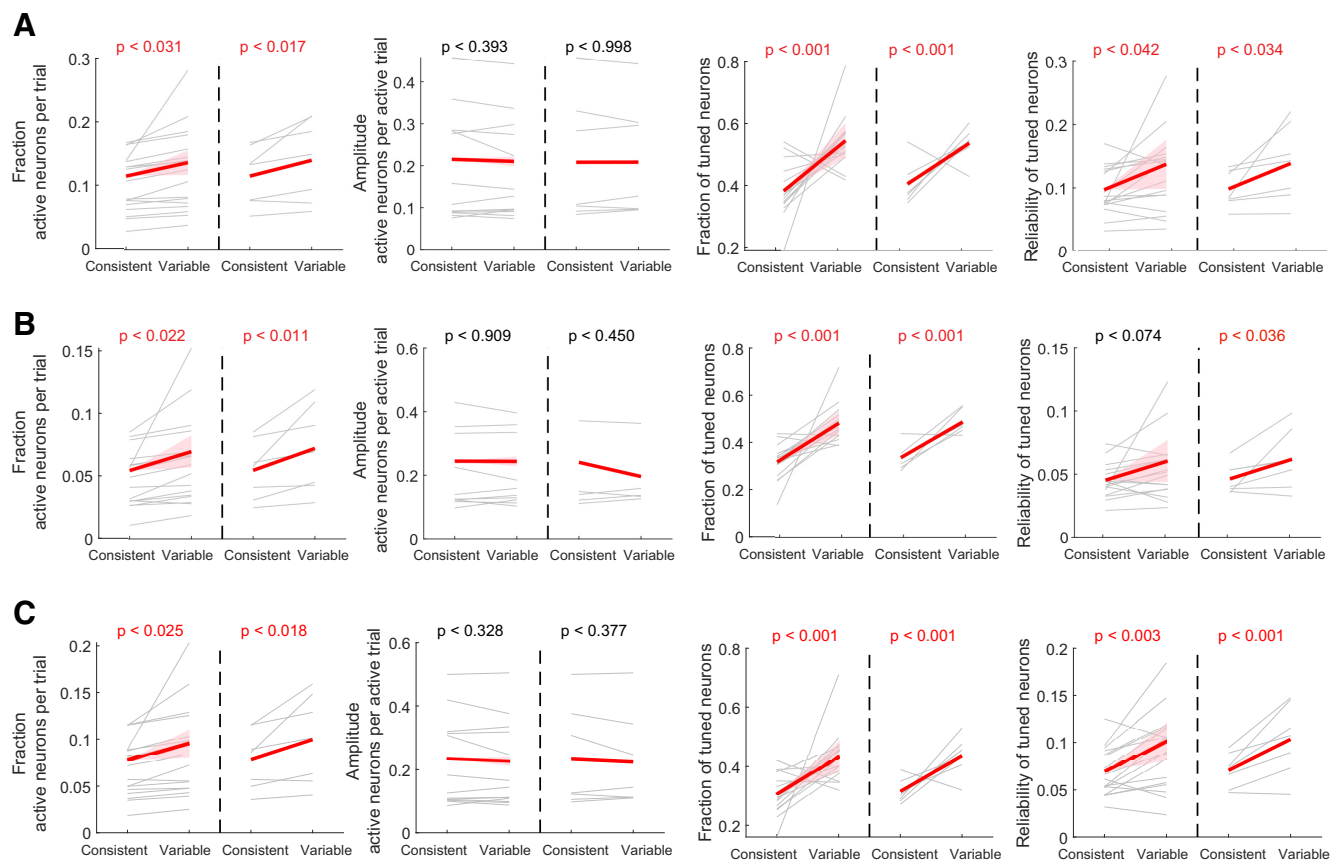
Figure 8. M1 activity comparison between the consistent- and variable-direction movements. *A*, The same format and analyses as Figure 7*F-I*, but using mean + 1.5 SD of the baseline activity as the per-neuron activation threshold. *B*, The same as in *A*, but using mean + 2.5 SD of the baseline activity as the per-neuron activation threshold. Note the similar trends in the direction dependency across different thresholds. *C*, The same as in *A*, but only using neurons that showed reliability difference between the two directions of at least 0.01. The trends are similar to the analyses using all neurons.
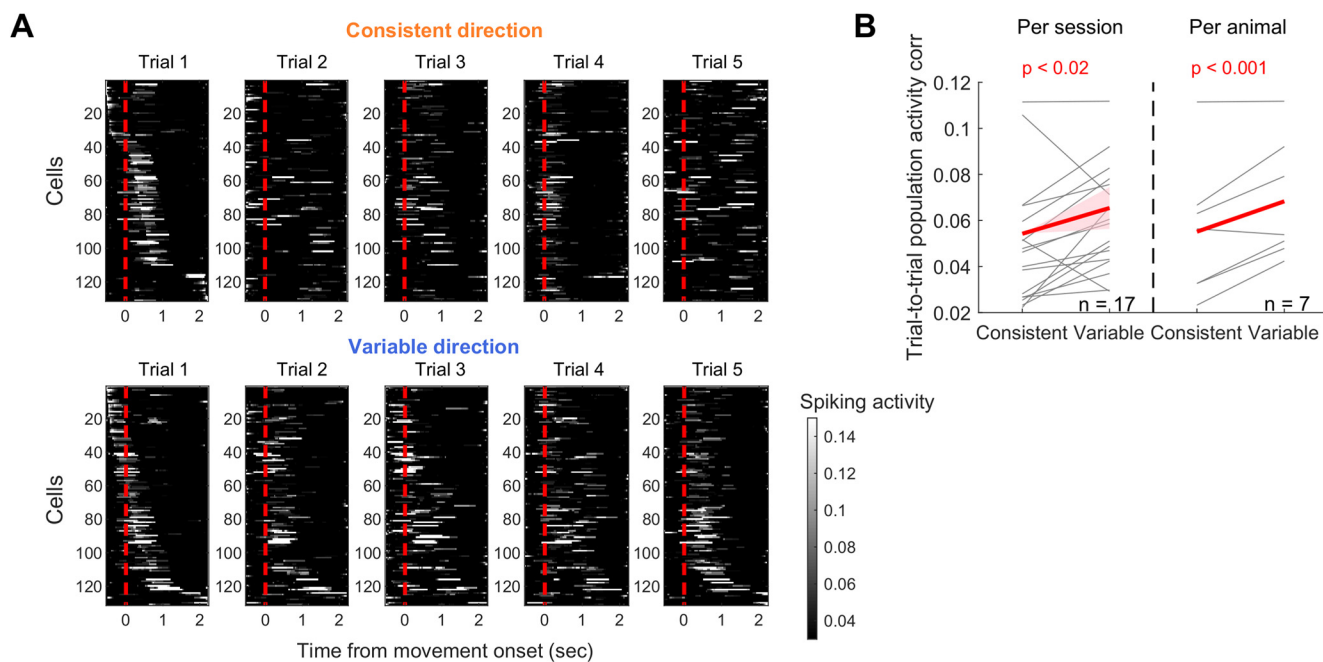
Figure 9. The consistency of M1 population activity modulated by movement consistency. *A*, M1 population activity during 5 rewarded trials in the consistent (top) versus variable (bottom) directions. Each row of the heat map represents a single neuron, and neurons are sorted based on the time of their peak activity in the first trial. *B*, The trial-to-trial population activity correlation for the consistent versus variable directions. The same illustration format and statistical tests as in Figure 7*C*.
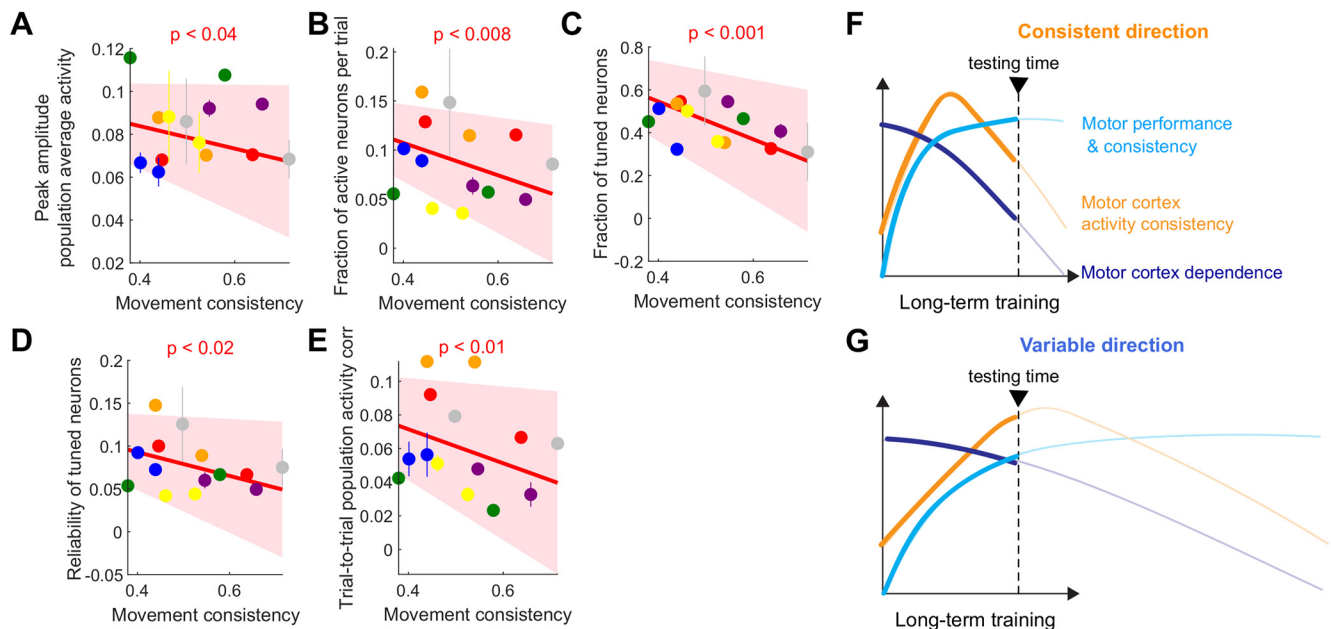
Figure 10. Disengagement of M1 modulated by motor performance. **A**, The peak amplitude of M1 population average activity as a function of movement consistency. Different colors for different mice. For illustration purpose, inactivation effect for each animal was shown as mean ± SE across imaging sessions. The dependence of the activity amplitude on the movement consistency was estimated using a linear mixed-effect model similar to Figure 6A. The significant slope indicates that the amplitude is significantly modulated by the movement consistency (Table 1). **B**, The fraction of M1 neurons activated per trial during perimovement period as a function of movement consistency. The same illustration and statistical test as in **A**. **C**, The fraction of tuned neurons as a function of movement consistency. The same illustration and statistical test as in **A**. **D**, The activation reliability of tuned neurons as a function of movement consistency. The same illustration and statistical test as in **A**. **E**, The trial-to-trial population activity correlation as a function of movement consistency. The same illustration and statistical test as in **A**. **F**, **G**, The change of motor performance and movement consistency, M1 activity consistency, and M1 dependence over the course of long-term training proposed by Hwang et al. (2019a). The two-directional movements in the current study take different time courses of learning. Movements in the consistent direction (top) may correspond to the later phase of learning than movements in the variable direction (bottom).

the two directions (d'Avella et al., 2006). The posture imposed by head fixation may make one direction more awkward to coordinate than the other. A combination of such factors that could differ from device to device and from animal to animal may underlie the different skill levels achieved between the two directions after a similar amount of training. A remaining open question is whether further training would allow animals to achieve a higher consistency for the variable-direction movement; thus, the variable-direction movement would become less dependent on M1, and its associated M1 activity would become less consistent. One possibility is that learning of distinct but similar movements may interfere with each other as they compete for common neural resources (Shadmehr and Brashers-Krug, 1997; Cantarero et al., 2013). In such a case, removing interference during further training by practicing the variable direction alone might result in improvement. Testing this possibility may further illuminate the neural mechanisms underlying motor learning of multiple novel movements and suggest efficient ways to teach complex motor skills consisting of various distinct movement components.

Increased incorrect choices in response to the variable-direction cue with M1 inactivation
In addition to the altered kinematics selectively in the variable-direction movement, we found that M1 inactivation increases the fraction of the consistent-direction movement in response to the cue instructing the variable-direction movement, but not vice versa. The increased rate of the consistent-direction movement, by itself, might suggest

various interpretations regarding the role of M1. First, M1 may perform sensorimotor association between visual cues and the appropriate motor goals, and M1 inactivation could impair this associative function, resulting in inaccurate target selection or inaccurate activation of a movement plan to the consistent direction, which may be the default preferred direction in our task. Second, M1 may be involved in suppressing highly consistent movements, in favor of less consistent movements to allow the exploration necessary for learning. Inactivation of M1 could then interfere with this selective suppression, increasing the rate of the consistent-direction movement.

Given other additional inactivation effects we observed, however, we favor the interpretation that M1 is actively involved in the control of variable-direction movements, whereas consistent-direction movements are controlled by M1-bypassing circuits. When animals choose a movement from a known, fixed set of alternatives as in our task, sensorimotor areas in the brain may prepare multiple movement plans until information to determine the correct response becomes available (Cisek and Kalaska, 2010; Klaes et al., 2011; Heindorf et al., 2018). In our task, the brain might simultaneously prepare both forward and downward reaches. Under the normal condition, the preparatory activity for the correct movement might prevail following the cue while the other degrades. However, when M1 is inactivated, the preparatory activity in M1 is perturbed, and the preparatory activity elsewhere may dominate the motor response. That is, the consistent movement prepared in the circuit bypassing M1 may not be suppressed when M1 is

inactivated, resulting in an increased fraction of trials in which the consistent-direction movement is executed. This interpretation fits well with our other findings that (1) even when variable-direction movements are made under M1 inactivation, they still exhibit altered kinematics; and (2) the normal M1 activity is strongly related to the kinematics of the variable-direction movement.

Benefits of M1 disengagement from learned movements

It is unknown why the control of learned movements moves from M1 to elsewhere. A tempting hypothesis is that the random explorative nature of M1 activity is adaptive for learning and refinement of new skills, but such an exploration may introduce unwanted variability for well-learned movements. So once movements are well learned and do not require further refinement any longer, it might be beneficial for the learned movement to be under the control of a neural pathway in which activity is more stable than M1. Analogous to this concept, it has been shown in songbirds that the relative influence of a songbird brain area LMAN, a source of motor variability that promotes song learning, gradually weakens while the influence of another brain area that generates stereotypic activity patterns increases as the learned song becomes crystalized (Garst-Orozco et al., 2014). Functional imaging and perturbation in the human also suggest that, following motor learning, the neural circuit representing the learned motor program changes and that this change may be associated with the increased functional stability of learned movements (Shadmehr and Holcomb, 1997; Muellbacher et al., 2002). Recent studies report that DLS might be an important node of an alternative pathway bypassing M1 in the late stage of learning (Dhawale et al., 2019; Wolff et al., 2019). Given the earlier evidence that automatic, habitual responses rely on DLS (Yin and Knowlton, 2006), a transfer of control from M1 to DLS might also facilitate the automatized execution of learned movements that is often a feature of highly skilled movements (Passingham, 1996; Schmidt et al., 2018). Such a transfer would also make the limited resources of M1 available for learning of new movements.

References

Berridge KC, Whishaw IQ (1992) Cortex, striatum and cerebellum: control of serial order in a grooming sequence. Exp Brain Res 90:275–290.

Biane JS, Takashima Y, Scanziani M, Conner JM, Tuszynski MH (2019) Reorganization of recurrent layer 5 corticospinal networks following adult motor training. J Neurosci 39:4684–4693.

Cantarero G, Tang B, O'Malley R, Salas R, Celnik P (2013) Motor learning interference is proportional to occlusion of LTP-like plasticity. J Neurosci 33:4634–4641.

Cisek P, Kalaska JF (2010) Neural mechanisms for interacting with a world full of action choices. Annu Rev Neurosci 33:269–298.

d'Avella A, Portone A, Fernandez L, Lacquaniti F (2006) Control of fast-reaching movements by muscle synergy combinations. J Neurosci 26:7791–7810.

Darling WG, Pizzimenti MA, Morecraft RJ (2011) Functional recovery following motor cortex lesions in non-human primates: experimental implications for human stroke patients. J Integr Neurosci 10:353–384.

Dhawale AK, Wolff SB, Ko R, Ölveczky BP (2019) The basal ganglia can control learned motor sequences independently of motor cortex. bioRxiv 827261.

Ebbesen CL, Brecht M (2017) Motor cortex: to act or not to act? Nat Rev Neurosci 18:694–705.

Garst-Orozco J, Babadi B, Ölveczky BP (2014) A neural circuit mechanism for regulating vocal variability during song learning in zebra finches. Elife 3:e03697.

Graziano M (2006) The organization of behavioral repertoire in motor cortex. Annu Rev Neurosci 29:105–134.

Guo JZ, Graves AR, Guo WW, Zheng J, Lee A, Rodríguez-Gonzá Lez J, Li N, Macklin JJ, Phillips JW, Mensh BD, Branson K, Hantman AW (2015) Cortex commands the performance of skilled movement. Elife 4:1–18.

Guo ZV, Li N, Huber D, Ophir E, Gutnisky D, Ting JT, Feng G, Svoboda K (2014) Flow of cortical activity underlying a tactile decision in mice. Neuron 81:179–194.

Heindorf M, Arber S, Keller GB (2018) Mouse motor cortex coordinates the behavioral response to unpredicted sensory feedback. Neuron 99:1040–1054.e5.

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

Hwang EJ, Dahlen JE, Hu YY, Aguilar K, Yu B, Mukundan M, Mitani A, Komiyama T (2019a) Disengagement of motor cortex from movement control during long-term learning. Sci Adv 5:eaay0001.

Hwang EJ, Link TD, Hu YY, Lu S, Wang EH, Lilascharoen V, Aronson S, O'Neil K, Lim BK, Komiyama T (2019b) Corticostriatal flow of action selection bias. Neuron 104:1126–1140.e6.

Kawai R, Markman T, Poddar R, Ko R, Fantana AL, Dhawale AK, Kampff AR, Ölveczky BP (2015) Motor cortex is required for learning but not for executing a motor skill. Neuron 86:800–812.

Klaes C, Westendorff S, Chakrabarti S, Gail A (2011) Choosing goals, not rules: deciding among rule-based action plans. Neuron 70:536–548.

Kupferschmidt DA, Juczewski K, Cui G, Johnson KA, Lovinger DM (2017) Parallel, but dissociable, processing in discrete corticostriatal inputs encodes skill learning. Neuron 96:476–489.e5.

Lemke SM, Ramanathan DS, Guo L, Won SJ, Ganguly K (2019) Emergent modular neural control drives coordinated motor actions. Nat Neurosci 22:1122–1131.

Lemon RN, Landau W, Tutssel D, Lawrence DG (2012) Lawrence and Kuypers (1968a, b) revisited: copies of the original filmed material from their classic papers in Brain. Brain 135:2290–2295.

Li N, Chen S, Guo ZV, Chen H, Huo Y, Inagaki HK, Chen G, Davis C, Hansel D, Guo C, Svoboda K (2019) Spatiotemporal constraints on optogenetic inactivation in cortical circuits. Elife 8:e48622.

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

Makino H, Hwang EJ, Hedrick NG, Komiyama T (2016) Circuit mechanisms of sensorimotor learning. Neuron 92:705–721.

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

Meamardoost S, Bhattacharya M, Hwang E, Komiyama T, Mewes C, Wang L, Zhang Y, Gunawan R (2020) FARCI: fast and robust connectome inference. bioRxiv 2020.10.07.330175.

Merel J, Botvinick M, Wayne G (2019) Hierarchical motor control in mammals and machines. Nat Commun 10:5489.

Muellbacher W, Ziemann U, Wissel J, Dang N, Kofler M, Facchini S, Boroojerdi B, Poewe W, Hallett M (2002) Early consolidation in human primary motor cortex. Nature 415:640–644.

Muir GD, Whishaw IQ (1999) Complete locomotor recovery following corticospinal tract lesions: measurement of ground reaction forces during overground locomotion in rats. Behav Brain Res 103:45–53.

Olsen SR, Bortone DS, Adesnik H, Scanziani M (2012) Gain control by layer six in cortical circuits of vision. Nature 483:47–52.

Owen SF, Liu MH, Kreitzer AC (2019) Thermal constraints on in vivo optogenetic manipulations. Nat Neurosci 22:1061–1065.

Papale AE, Hooks BM (2018) Circuit changes in motor cortex during motor skill learning. Neuroscience 368:283–297.

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

Peters AJ, Chen SX, Komiyama T (2014) Emergence of reproducible spatiotemporal activity during motor learning. Nature 510:263–267.

Peters AJ, Lee J, Hedrick NG, O'Neil K, Komiyama T (2017a) Reorganization of corticospinal output during motor learning. Nat Neurosci 20:1133–1141.

Peters AJ, Liu H, Komiyama T (2017b) Learning in the rodent motor cortex. Annu Rev Neurosci 40:77–97.

Sanes JN, Donoghue JP (2000) Plasticity and primary motor cortex. Annu Rev Neurosci 23:393–415.

Sauerbrei BA, Guo JZ, Cohen JD, Mischiati M, Guo W, Kabra M, Verma N, Mensh B, Branson K, Hantman AW (2020) Cortical pattern generation during dexterous movement is input-driven. Nature 577:386–391.

Schmidt RA, Lee TD, Winstein CJ, Wulf G, Zelaznik H (2018) Motor control and learning, Ed 6. Champaign, IL: Human Kinetics.

Shadmehr R, Brashers-Krug T (1997) Functional stages in the formation of human long-term motor memory. J Neurosci 17:409–419.

Shadmehr R, Holcomb HH (1997) Neural correlates of motor memory consolidation. Science 277:821–825.

Ungerleider LG, Doyon J, Karni A (2002) Imaging brain plasticity during motor skill learning. Neurobiol Learn Mem 78:553–564.

Wolff SB, Ko R, Ölveczky BP (2019) Distinct roles for motor cortical and thalamic inputs to striatum during motor learning and execution. bioRxiv 825810.

Yin HH, Knowlton BJ (2006) The role of the basal ganglia in habit formation. Nat Rev Neurosci 7:464–476.