LETTER

doi:10.1038/nature10845

Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills

Aaron C. Koralek[1]*, Xin Jin[5]*, John D. Long II[1], Rui M. Costa[5,6] & Jose M. Carmena[1,2,3,4]

The ability to learn new skills and perfect them with practice applies not only to physical skills but also to abstract skills[1], like motor planning or neuroprosthetic actions. Although plasticity in corticostriatal circuits has been implicated in learning physical skills[2–4], it remains unclear if similar circuits or processes are required for abstract skill learning. Here we use a novel behavioural task in rodents to investigate the role of corticostriatal plasticity in abstract skill learning. Rodents learned to control the pitch of an auditory cursor to reach one of two targets by modulating activity in primary motor cortex irrespective of physical movement. Degradation of the relation between action and outcome, as well as sensory-specific devaluation and omission tests, demonstrate that these learned neuroprosthetic actions are intentional and goal-directed, rather than habitual. Striatal neurons change their activity with learning, with more neurons modulating their activity in relation to target-reaching as learning progresses. Concomitantly, strong relations between the activity of neurons in motor cortex and the striatum emerge. Specific deletion of striatal NMDA receptors impairs the development of this corticostriatal plasticity, and disrupts the ability to learn neuroprosthetic skills. These results suggest that corticostriatal plasticity is necessary for abstract skill learning, and that neuroprosthetic movements capitalize on the neural circuitry involved in natural motor learning.

The ability to learn new actions and perfect them with practice allows us to master skills like playing the piano or riding a bicycle. Learning these skills usually implies moving faster, more accurately and less variably[5]. However, mastering other types of skills, like playing board games or controlling neuroprosthetic devices, often does not directly involve changes in physical movement[1,6]. Cortico-basal ganglia circuits have been implicated in the learning, selection and execution of physical skills[2–4,7,8]. In particular, plasticity in the motor cortices and the striatum, the major input region of the basal ganglia, has been shown to accompany the learning of physical skills[2,9]. The motor cortex and frontal cortices have also been implicated in the learning of abstract skills[10–13], and in learning to control neuroprosthetic devices irrespective of physical movement[14–17]. Some studies suggest that not only cortical areas, but also the striatum, are involved in learning abstract skills[18–20]. However, it is still unclear if the striatum is required for abstract skill learning, and if corticostriatal circuits undergo plasticity during the learning of such skills as they do during the learning of physical skills. Here, we use a novel behavioural task in conjunction with electrophysiology and genetic manipulation in rodents to investigate the role of corticostriatal circuits and corticostriatal plasticity in the learning of intentional neuroprosthetic actions: that is, actions performed with disembodied actuators based on the modulation of specific neural activity and irrespective of physical movement[6].

We developed a novel operant brain–machine interface task in which rodents were required to modulate activity in M1, rather than execute a physical movement, to obtain reward (Fig. 1a). Modulation of M1 ensemble activity resulted in changes in the pitch of an auditory cursor, which provided constant auditory feedback to rodents about task performance. Reward was delivered when rodents precisely modulated M1 activity to move this auditory cursor to one of two

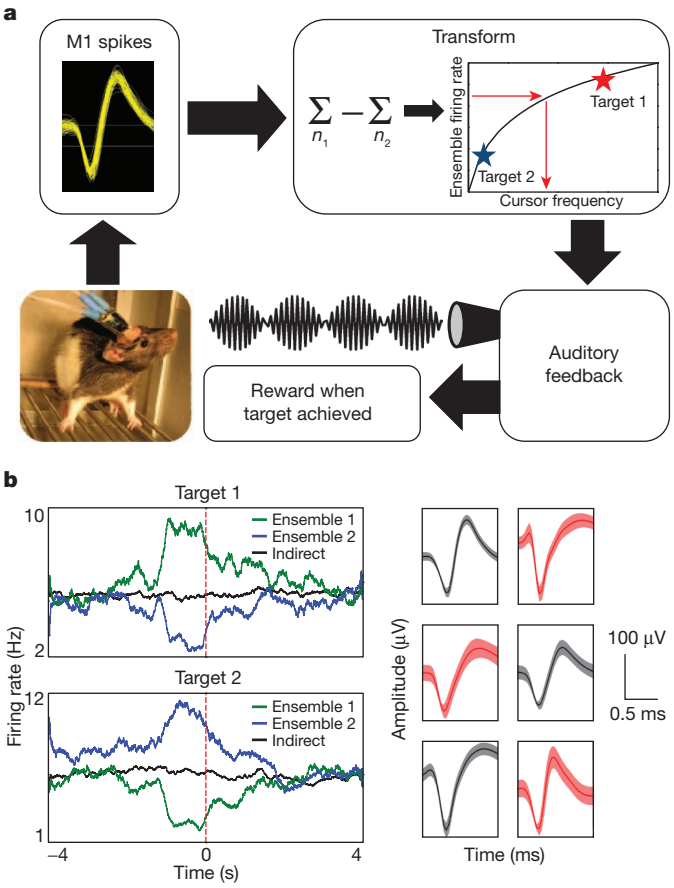

Figure 1 | Volitional modulation of M1 neural activity in awake behaving rodents. **a**, Task schematic. M1 unit activity was entered into an online transform algorithm that related ensemble activity to the pitch of an auditory cursor. Two opposing ensembles were chosen, with activity of one ensemble increasing the cursor pitch and activity of the other ensemble decreasing the cursor pitch. Constant auditory feedback about cursor location was supplied to rodents, and distinct rewards were supplied when rodents brought M1 activity into one of two target states. **b**, Mean M1 ensemble firing rates for units in ensemble 1 (green), ensemble 2 (blue) and M1 units not used in the transform (black) in relation to the achievement of target 1 (top) or target 2 (bottom). Representative waveforms recorded from M1 in rats are shown on the right, with shaded regions denoting the standard deviation.

[1]Helen Wills Neuroscience Institute, University of California, Berkeley, California 94720, USA. [2]Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, California 94720, USA. [3]Program in Cognitive Science, University of California, Berkeley, California 94720, USA. [4]UC Berkeley and UC San Francisco Joint Graduate Group in Bioengineering, University of California, Berkeley, California 94720, USA. [5]Laboratory for Integrative Neuroscience, National Institute on Alcohol Abuse and Alcoholism, National Institutes of Health, 5625 Fishers Lane, Bethesda, Maryland 20892-9412, USA. [6]Champalimaud Neuroscience Programme, Champalimaud Center for the Unknown, Avenida de Brasília, 1400-038 Lisbon, Portugal.
*These authors contributed equally to this work.

target tones, and a trial was marked incorrect if no target had been hit within a set time limit (30 s). One of these targets was associated with a reward of sucrose solution, whereas the other target was associated with a pellet reward (see Methods). Two neural ensembles consisting of two- to four-well-isolated units each were used to control the auditory cursor (Supplementary Figs 1 and 2). The action of these two ensembles opposed each other, such that increased activity in one ensemble produced increases in cursor pitch, whereas increased activity in the other ensemble caused decreases in cursor pitch. Thus, to achieve a high-pitched target, rodents had to increase activity in the first ensemble and decrease activity in the second; the opposite was required to hit a low-pitched target (Fig. 1b and Supplementary Fig. 3). These firing-rate modulations had to be maintained for several time bins (200 ms bin size) for a target to be hit (Supplementary Methods). Hence, in this operant task, rodents had to bring the two M1 ensembles into a desired state irrespective of motor output.

We trained six male Long–Evans rats on the task, and verified that they exhibited marked improvement in the percentage of correct trials over the course of 11 days (Fig. 2a). As typically observed in motor skill learning[21], there was a phase of rapid improvement followed by a phase of slower learning, representing early (days 2–4) and late (days 8–11) phases of learning. The percentage of correct trials increased significantly from early to late in learning (Fig. 2b; $P < 0.001$), resulting in performance well above chance (Fig. 2c; $P < 0.001$; see Supplementary Methods), whereas mean time-to-target decreased (Supplementary Fig. 4). Analyses of M1 firing rates further showed that rats were producing the desired neuronal ensemble rate modulations during task performance (Fig. 1b and Supplementary Fig. 3). Furthermore, sensory feedback was found to be critical for animals to learn this task because when rats were not given auditory feedback during training (although they would still get a reward if they would modulate neural activity correctly), the percentage of correct trials did not increase over the course of 9 days of training (Supplementary Fig. 5).

We next investigated if animals were performing physical movements that would modulate the activity of those particular M1 ensembles. First, we monitored overall rodent movement with an accelerometer mounted on the recording headstage, which allowed us to measure if the animals produced any body or head movement during target achievement[22]. Accelerometer traces exhibited no changes before and during target reaching, but did show prominent deflections after target reaching as the animals retrieved the reward (Fig. 2d), demonstrating that rodents were not relying on gross motor behaviour to perform the task. We also monitored movements of the vibrissae with electromyographic (EMG) recordings of the mystacial pad (electrodes targeted M1 areas controlling vibrissae movement; Supplementary Fig. 2 and Supplementary Methods), and observed no significant EMG signals before target achievement, although there were clear EMG signals afterwards as animals retrieved and consumed the reward (Fig. 2e and Supplementary Fig. 6b). Importantly, there was no correlation between EMG activity and the spiking of the M1 neurons controlling the auditory cursor: the correlation coefficient for all trials in a behavioural session was 0.092 ± 0.003 (mean $\pm$ s.e.m.), and the distribution of correlation coefficients across a session was not significantly different from zero (Supplementary Fig. 6a; $P = 0.57$). This was observed across all training days, including during early learning (Supplementary Fig. 7). These data suggest that rats do not rely on physical movements to learn the task, although it is difficult to exclude the possibility that animals use some movement to generate neural activity to drive the auditory cursor during exploratory phases of the task in early learning. Nonetheless, the data show that animals eventually learn to perform the task in the absence of overt movement. To demonstrate further that rats did not require vibrissae movements to control M1 activity, we injected lidocaine into the whisker pad to inactivate sensory and motor nerve endings locally during a session in late learning (see Supplementary Methods). There was no significant change in performance during the temporary inactivation (Fig. 2f;

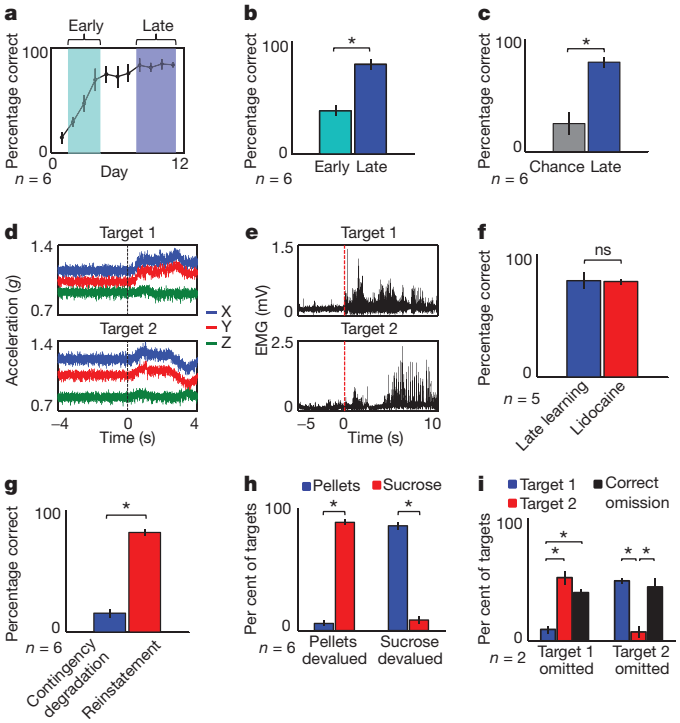

Figure 2 | Learning intentional neuroprosthetic actions independently of movement. **a**, Mean percentage of correct responses for all rats across days 1–11 of learning. Shaded regions denote the range of days from which the early and late learning analyses were performed. **b**, The percentage of correct responses for all rats increased significantly from early (light blue) to late (dark blue) in learning. **c**, Percentage of correct responses in late learning is significantly greater than expected by chance. **d**, Representative accelerometer traces show no gross motor behaviour leading to target achievement (time zero), but clear deflections as animals initiate movement to retrieve reward. **e**, Representative EMG traces show no muscle activity in the mystacial pad before target achievement, but clear deflections as animals retrieve and consume reward. **f**, Mean performance in all rats when lidocaine was injected into the whisker pad before a behavioural session late in learning (red) compared with performance during a no-lidocaine session (dark blue). **g**, Significant reduction in response rate when the causal relation between target achievement and reward delivery was degraded (dark blue). When contingency was reinstated, performance rapidly returned to pre-degradation levels (red). **h**, Percentage of total correct trials that were directed at the target associated with pellet reward (blue) or sucrose solution reward (red) during choice sessions where rats had free access to pellets (left; 'Pellets devalued'), or to sucrose before the session (right; 'Sucrose devalued'). **i**, Percentage of total trials that involved responses towards target 1 (blue), target 2 (red) or response omissions (black) when omission tests were performed for target 1 (left) or target 2 (right).

$P > 0.9$), with rats achieving $78.1 \pm 2.2\%$ correct with lidocaine (mean $\pm$ s.e.m.) versus $78.8 \pm 6.5\%$ without lidocaine. Taken together, these data indicated that rodents were able to learn to control M1 activity operantly irrespective of any overt movement.

Goal-directed actions are sensitive to changes in the relation between performing the action and obtaining a reward (contingency) and to changes in the expected value of the reward, whereas habits are not[23,24]. We asked if these neuroprosthetic actions were performed intentionally, because the animal volitionally controlled M1 activity to get the outcome (goal-directed), or habitually owing to the reinforcement history. To test this, we first degraded the contingency between executing the action and obtaining the outcome: that is, the auditory cursor was still under control of M1 ensemble activity, but the probability of obtaining reward was similar irrespective of target achievement, which had no effect on the rate of reward. After 2 days of contingency degradation, rats markedly diminished their responding and the percentage of correct trials decreased significantly (Fig. 2g;

$P < 0.001$). When contingency was reinstated, rats resumed responding and the percentage of correct trials returned to plateau levels seen in late learning (Fig. 2g).

To investigate further the intentional nature of these neuroprosthetic skills, we performed a test where each of the outcomes was devalued using sensory-specific satiety. Rats were given free access either to sucrose solution or pellets for 1 h before the behavioural session, thereby reducing the expected value of that outcome[25]. After specific devaluation of each outcome/reward, rats chose the target leading to that reward much less than the target leading to the reward that was not devalued (Fig. 2h; $P < 0.001$), indicating that their actions were sensitive to changes in outcome value. Importantly, there were no significant differences in reward preference during normal task performance when neither of the outcomes was devalued (Supplementary Fig. 8; $P > 0.25$). Finally, we asked whether rats were able to intentionally inhibit the reaching of one of the two targets to obtain the specific reward associated with that target. To examine this we performed an omission test, where the reward previously associated with reaching a particular target was only delivered when rats successfully inhibited reaching that target throughout the duration of the trial. If the target was reached during the 30 s of trial duration, no reward was delivered and a new trial was initiated. Importantly, reaching the other target continued to lead to reward as during training. Animals behaved in a goal-directed manner in the omission test for both targets, because they reduced the number of reaches for the target they had to omit versus the no-omission target, while increasing the number of correctly omitted responses (Fig. 2i; $P < 0.001$ for both comparisons). Taken together, these data show that the neuroprosthetic actions in our task are sensitive to changes in the causal relation between performing the action and obtaining the reward (contingency degradation and omission test), and to changes in the expected value of the outcome (sensory-specific devaluation), indicating that they are intentional and goal-directed rather than habitual.

We next examined if learning to operantly control M1 activity irrespective of overt movement involves striatal plasticity, akin to what is observed for natural motor learning[2–4,7,26–28]. We verified that the improvement in behavioural performance seen across learning was accompanied by a significant increase in firing rates in the dorsal striatum (DS) in late learning compared with early learning (Fig. 3a; $P < 0.001$). In addition to this general increase in firing rates, we noticed that firing rates of DS neurons exhibited greatest modulation during target reaching compared with baseline control periods (Fig. 3b), as observed during natural motor learning[26]. This modulation was significantly greater in late learning than early learning (Fig. 3b, c; $P < 0.05$), indicating that DS neurons changed their activity during the volitional control of M1 activity, and that this change increased with learning.

We next investigated if learning, and the observed changes in DS target-related activity, were accompanied by corticostriatal plasticity, that is, changes in the functional interactions between M1 and DS neurons. We noticed that cross-correlation histograms between the two regions in late learning exhibited pronounced oscillatory spike coupling (Fig. 3d). To quantify this interaction, we calculated the coherence between spiking activity in the two regions in both early and late learning (Supplementary Methods). The resulting coherograms exhibited a clear increase in coherence at low-frequency bands in late learning relative to early learning (Fig. 3e), and these frequencies corresponded to the oscillatory frequency seen in the cross-correlograms (Supplementary Methods). Furthermore, mean coherence in the theta band (4–8 Hz) was significantly greater in late learning than early learning (Fig. 3f; $P < 0.001$). This increase in coherence appeared to be related to learning to perform the task rather than higher reward expectation or proportion of correct trials in late learning, because coherence values remained high surrounding target achievement during the contingency degradation manipulation, where reward delivery was not contingent upon target achievement (Supplementary

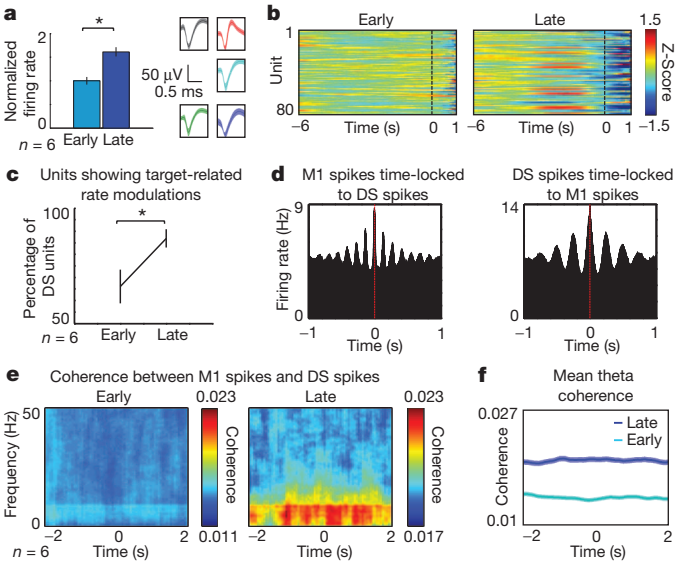

Figure 3 | Learning abstract skills is accompanied by corticostriatal plasticity. **a**, Mean normalized firing rates in DS increased significantly from early (light blue) to late (dark blue) learning. Representative waveforms recorded from the DS are shown on the right (shaded regions denote standard deviation). **b**, *Z*-scored firing rates for individual DS units in relation to target achievement (time zero) showing marked modulation before target achievement in late learning. **c**, The percentage of DS units exhibiting target-related firing-rate modulation increased significantly with learning. **d**, Cross-correlation histograms in late learning for M1 spiking activity in relation to DS spikes (left), and DS spiking activity in relation to M1 spikes (right), showing oscillatory coupling between the two regions. **e**, Coherence between M1 spikes and DS spikes in early (left) and late (right) learning shows a clear increase in low-frequency coherence from early to late learning. **f**, Significant increase in mean coherence in the theta range in late (dark blue) versus early (light blue) learning. Shaded regions denote s.e.m.

Fig. 9; not significantly less than non-degraded trials, $P > 0.05$). In addition, coherence levels remained high during task performance in incorrect trials (data not shown), further suggesting that the increase in coherence observed is due to learning to perform the skill rather than outcome anticipation. Thus, neuroprosthetic skill learning is accompanied by dynamic changes in functional interactions between M1 and the DS neurons, suggesting an important role for corticostriatal plasticity in this novel task.

We therefore investigated if corticostriatal plasticity would be necessary for neuroprosthetic skill learning. NMDARs (*N*-methyl-D-aspartic acid receptors) in striatal medium spiny neurons are critical for corticostriatal long-term potentiation[29]. We used a knockin line that expresses Cre recombinase in both striatonigral and striatopallidal medium spiny neurons (*RGS9L-cre*), but not in all striatal neurons (for example, absent from parvalbumin interneurons; Supplementary Fig. 10), and crossed it with mice carrying a floxed allele of the *NMDAR1* gene[30]. The resulting mice lack NMDA currents in most projection neurons[30] (but not all striatal cells, hence we refer to them as *RGS9L-Cre/Nr1$^{f/f}$*, not as striatal NR1 knockouts: see Supplementary Methods), and have impaired corticostriatal long-term potentiation[30]. As previously described, these animals do not display any major motor deficits (Supplementary Videos 1 and 2) and can learn to perform rapid sequential movements (Supplementary Fig. 11), albeit being unable to learn precise motor sequences[27]. We investigated neuroprosthetic skill learning in *RGS9L-Cre/Nr1$^{f/f}$* mice and littermate controls. Although control mice showed performance improvement across learning irrespective of physical movement as observed for rats ($P < 0.001$; Fig. 4a, b), *RGS9L-Cre/Nr1$^{f/f}$* mice exhibited marked learning deficits on the task, with no significant increase in the percentage of correct trials from early to late learning (Fig. 4a; $P = 0.98$). Furthermore, acute pharmacological blockade of NMDARs in trained

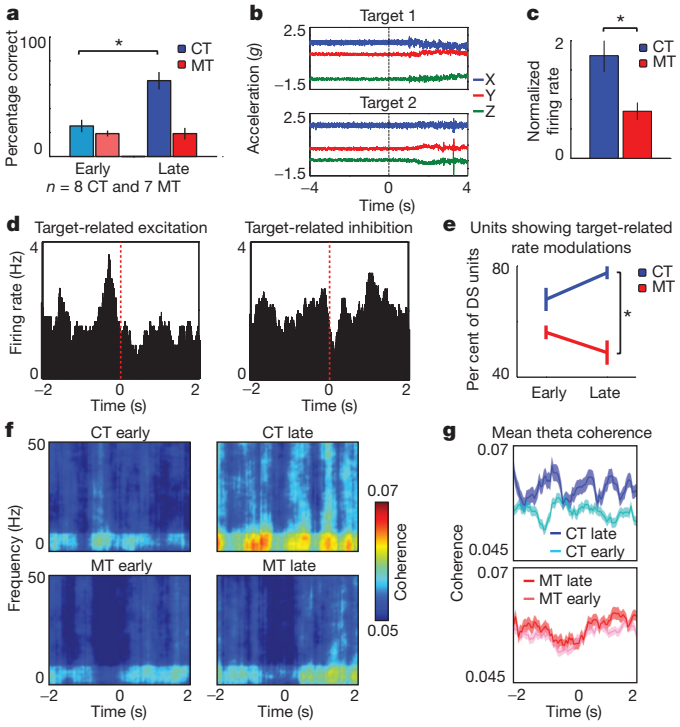

Figure 4 | Selective deletion of NMDARs in the striatum impairs brain–machine interface learning. **a**, *RGS9L-Cre/Nr1$^{f/f}$* mice (red) exhibit no significant increase in the percentage of correct trials over the course of learning, despite clear performance improvement in littermate controls (blue). CT, controls; MT, mutants. **b**, Accelerometer traces from control mice showing no clear oscillation before target achievement, but clear deflections as mice retrieve reward. **c**, Late firing rate normalized to early firing rate in controls and mutants. There is no significant increase in DS firing rates in *RGS9L-Cre/Nr1$^{f/f}$* mice (red) from early to late learning, although DS firing rates increase markedly in control mice (blue). **d**, DS units of controls exhibit strong target-related firing-rate modulations, including both excitation (left) and inhibition (right). **e**, The percentage of DS units showing significant target-related firing-rate modulations increases significantly across learning in control mice (blue), but not in *RGS9L-Cre/Nr1$^{f/f}$* mutants (red). **f**, Coherograms showing coherence between M1 spikes and DS spikes in early (left) and late (right) learning both for control mice (top) or *RGS9L-Cre/Nr1$^{f/f}$* mice (bottom). Coherence in low-frequency bands increases from early to late learning in control mice, but not in *RGS9L-Cre/Nr1$^{f/f}$* mice. **g**, Mean coherence in the theta range for control (top) and mutant (bottom) mice. There is a significant increase in coherence from early (light blue) to late (dark blue) learning in control mice but not in mutant mice (early learning, light red; late learning, dark red).

control animals did not affect performance of the neuroprosthetic skill (even at relatively high doses that affect striatal burst firing; Supplementary Figs 12 and 13), suggesting that the deficits observed in *RGS9L-Cre/Nr1$^{f/f}$* mice are not due to inability to perform the skill but rather to the inability to learn the task.

Consistent with the findings above in rats (Fig. 3a), DS neurons in littermate control mice exhibited a significant increase in firing rate across learning, whereas in mutants they did not (Fig. 4c; main effect of genotype $F_{1,10} = 32.45$, $P < 0.001$; early versus late $P < 0.05$ for controls (CT) and $P = 0.23$ for mutants (MT)). Also, in control mice, the proportion of DS neurons with significant target-related firing-rate modulation increased with learning (Fig. 4d, e; $P < 0.05$), but this was not observed in *RGS9L-Cre/Nr1$^{f/f}$* mice (Fig. 4e; $P = 0.28$). Finally, the development of functional corticostriatal interactions during learning was also abolished in *RGS9L-Cre/Nr1$^{f/f}$* mutants, with no significant increase in coherence between M1 and DS spikes with learning (Fig. 4f, g, $F_{80,10} = 0.65$, $P = 0.44$), although littermate controls showed a clear increase as seen in rats (Fig. 4f, g; $F_{80,10} = 4.86$, $P < 0.05$). Taken together, these results demonstrate that the striking corticostriatal plasticity observed in rats during learning also occurs in control mice, but this plasticity is absent in mice with a decrease in functional NMDARs in striatum. These mutant mice do not show improvement with training, therefore indicating that corticostriatal plasticity may be necessary for learning to modulate M1 states intentionally to obtain specific outcomes.

In summary, we used a novel operant task in rodents to demonstrate that corticostriatal networks exhibit profound plasticity during the learning of intentional neuroprosthetic skills and, further, that disrupting this plasticity impairs learning. This adds great support to the claims that cortico-basal ganglia circuits play a role in abstract cognitive processes[18–20]. We observed that DS neurons strongly modulated their activity in relation to M1 activity, even when the latter was dissociated from physical movements, suggesting that the striatum is important for learning and selecting abstract actions that are controlled by cortical output. Hence, these data suggest that cortico-basal ganglia circuits may be involved in learning mental actions and skills that do not require physical movement, indicating that they may have a broader function involved in intention and decision-making than previously acknowledged.

Our results also have important implications for the field of brain–machine interfaces[6]. The abstract actions investigated here form the basis for skilful neuroprosthetic control[16] and, as we have shown here, they recruit elements of the natural motor system outside of M1. Thus, our results suggest that neuroprosthetic movements capitalize on the neural circuitry for motor learning and therefore have great potential to feel naturalistic, generalize well to novel movements and environments, and benefit from our nervous system's highly developed storage and retrieval mechanisms for skilled behaviour.

METHODS SUMMARY

All experiments were done in accordance with the Animal Care and Use Committee at the University of California, Berkeley, and at the National Institute on Alcohol Abuse and Alcoholism, and according to National Institutes of Health guidelines. Six male Long–Evans rats, seven *RGS9L-Cre/Nr1$^{f/f}$* mice and eight littermate controls were chronically implanted with tungsten microelectrode arrays in both M1 and the DS ipsilaterally. Two ensembles of two to four well-isolated M1 units each were chosen and, by modulating activity in these ensembles, rodents controlled the pitch of an auditory cursor, with increased activity in the first ensemble producing increases in the cursor pitch and increased activity in the other ensemble producing decreases in the cursor pitch. The rodents had to modulate these ensembles precisely to move the cursor to one of two target pitches to get reward (one associated with 20% sucrose and another with pellets). Rodents were free to choose either reward at any time. A trial was correct if a target was achieved within 30 s and incorrect trials were followed by a time-out. M1 activity levels had to return to baseline levels for a new trial to begin. After performance had reached plateau levels, the action–outcome contingency was degraded by providing outcomes on a variable time schedule to match the probabilities between getting a reward after target achievement and no target achievement. For outcome devaluation, rodents were given a sensory-specific satiety test where they received free access to one of the rewards (but not the other) for 1 h before a behavioural session. For the omission test, rodents stopped being rewarded for reaching one of the targets and the reward associated with that target was instead supplied when rodents successfully inhibited responses for the duration of the trial. See Supplementary Information for further details.

Received 21 April 2011; accepted 9 January 2012.
Published online 4 March 2012.

1. VanLehn, K. Cognitive skill acquisition. *Annu. Rev. Neurosci.* **47,** 513–539 (1996).
2. Yin, H. H. *et al.* Dynamic reorganization of striatal circuits during the acquisition and consolidation of a skill. *Nature Neurosci.* **12,** 333–341 (2009).
3. Barnes, T. D., Kubota, Y., Hu, D., Jin, D. Z. & Graybiel, A. M. Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. *Nature* **437,** 1158–1161 (2005).
4. Kimchi, E. Y. & Laubach, M. Dynamic encoding of action selection by the medial striatum. *J. Neurosci.* **29,** 3148–3159 (2009).
5. Brashers-Krug, T., Shadmehr, R. & Bizzi, E. Consolidation in human motor memory. *Nature* **382,** 252–255 (1996).
6. Fetz, E. E. Volitional control of neural activity: implications for brain–computer interfaces. *J. Physiol. (Lond.)* **579,** 571–579 (2007).
7. Hikosaka, O. *et al.* Parallel neural networks for learning sequential procedures. *Trends Neurosci.* **22,** 464–471 (1999).

8. Brasted, P. J. & Wise, S. P. Comparison of learning-related neuronal activity in the dorsal premotor cortex and striatum. *Eur. J. Neurosci.* **19,** 721–740 (2004).
9. Rioult-Pedotti, M. S., Friedman, D. & Donghue, J. P. Learning-induced LTP in neocortex. *Science* **290,** 533–536 (2000).
10. Georgopoulos, A. P., Taira, M. & Lukashin, A. Cognitive neurophysiology of the motor cortex. *Science* **260,** 47–52 (1993).
11. Gandolfo, F., Li, C., Benda, B. J., Schioppa, C. P. & Bizzi, E. Cortical correlates of learning in monkeys adapting to a new dynamical environment. *Proc. Natl Acad. Sci. USA* **97,** 2259–2263 (2000).
12. Fincham, J. M. & Anderson, J. R. Distinct roles of the anterior cingulate and prefrontal cortex in the acquisition and performance of a cognitive skill. *Proc. Natl Acad. Sci. USA* **103,** 12941–12946 (2006).
13. Badre, D., Kayser, A. S. & D'Esposito, M. Frontal cortex and the discovery of abstract action rules. *Neuron* **66,** 315–326 (2010).
14. Taylor, D. M., Tillery, S. I. & Schwartz, A. B. Direct cortical control of 3D neuroprosthetic devices. *Science* **296,** 1829–1832 (2002).
15. Carmena, J. M. *et al.* Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biol.* **1,** 193–208 (2003).
16. Ganguly, K. & Carmena, J. M. Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biol.* **7,** e1000153 (2009).
17. Ganguly, K., Dimitrov, D. F., Wallis, J. D. & Carmena, J. M. Reversible large-scale modification of cortical networks during neuroprosthetic control. *Nature Neurosci.* **14,** 662–667 (2011).
18. Beauchamp, M. H., Dagher, A., Aston, J. A. & Doyon, J. Dynamic functional changes associated with cognitive skill learning of an adapted version of the Tower of London task. *Neuroimage* **20,** 1649–1660 (2003).
19. Poldrack, R. A., Prabhakaran, V., Seger, C. A. & Gabrieli, J. D. Striatal activation during acquisition of a cognitive skill. *Neuropsychology* **13,** 564–574 (1999).
20. Pasupathy, A. & Miller, E. K. Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature* **433,** 873–876 (2005).
21. Karni, A. *et al.* The acquisition of skilled motor performance: fast and slow experience-driven changes in primary motor cortex. *Proc. Natl Acad. Sci. USA* **95,** 861–868 (1998).
22. Venkatraman, S., Jin, X., Costa, R. M. & Carmena, J. M. Investigating neural correlates of behavior in freely behaving rodents using inertial sensors. *J. Neurophysiol.* **104,** 569–575 (2010).
23. Balleine, B. W. & Dickinson, A. Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology* **37,** 407–419 (1998).
24. Yin, H. H., Knowlton, B. J. & Balleine, B. W. Inactivation of dorsolateral striatum enhances sensitivity to changes in the action-outcome contingency in instrumental conditioning. *Behav. Brain Res.* **166,** 189–196 (2006).
25. Hilario, M. R., Clouse, E., Yin, H. H. & Costa, R. M. Endocannabinoid signaling is critical for habit formation. *Front. Integr. Neurosci.* **1,** 1–12 (2007).
26. Costa, R. M., Cohen, D. & Nicolelis, M. A. Differential corticostriatal plasticity during fast and slow motor skill learning in mice. *Curr. Biol.* **14,** 1124–1134 (2004).
27. Jin, X. & Costa, R. M. Start/stop signals emerge in nigrostriatal circuits during sequence learning. *Nature* **466,** 457–462 (2010).
28. Miyachi, S., Hikosaka, O. & Lu, X. Differential activation of monkey striatal neurons in the early and late stages of procedural learning. *Exp. Brain Res.* **146,** 122–126 (2002).
29. Calabresi, P., Pisani, A., Mercuri, N. B. & Bernardi, G. Long-term potentiation in the striatum is unmasked by removing the voltage-dependent magnesium block of NMDA receptor channels. *Eur. J. Neurosci.* **4,** 929–935 (1992).
30. Dang, M. T. *et al.* Disrupted motor learning and long-term synaptic plasticity in mice lacking NMDAR1 in the striatum. *Proc. Natl Acad. Sci. USA* **103,** 15254–15259 (2006).

Supplementary Information is linked to the online version of the paper at www.nature.com/nature.

Acknowledgements We thank S. Venkatraman for the three-axis accelerometer, Y. Li for the *RGS9L-Cre* mice, K. Nakazawa for the *NMDAR1-loxP* mice, G. Luo for genotyping, M. Davis for advice on staining and G. Martins for performing immunohistochemistry. This work was supported by National Science Foundation CAREER Award 0954243, the Multiscale Systems Research Center and the Defense Advanced Research Projects Agency contract N66001-10-C-2008 to J.M.C., and the Division of Intramural Clinical and Basic Research of the National Institute on Alcohol Abuse and Alcoholism, Marie Curie International Reintegration Grant 239527 and European Research Council STG 243393 to R.M.C.

Author Contributions A.C.K., X.J., J.D.L., R.M.C. and J.M.C. designed experiments. A.C.K., X.J. and J.D.L. conducted experiments. A.C.K., X.J., R.M.C. and J.M.C. analysed data and wrote the paper.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of this article at www.nature.com/nature. Correspondence and requests for materials should be addressed to J.M.C. (carmena@eecs.berkeley.edu) or R.M.C. (ruicosta@fchampalimaud.org).