BRAIN
ORIGINAL ARTICLE

Volitional control of individual neurons in the human brain

Kramay Patel,[1,2,3,4] Chaim N. Katz,[1,2,4] Suneil K. Kalia,[1,4,5] Milos R. Popovic[1,2,4,6] and Taufik A. Valiante[1,2,4,5,6,7,8]

See Rubin and Paulk (doi:10.1093/brain/awab413) for a scientific commentary on this article.

Brain–machine interfaces allow neuroscientists to causally link specific neural activity patterns to a particular behaviour. Thus, in addition to their current clinical applications, brain–machine interfaces can also be used as a tool to investigate neural mechanisms of learning and plasticity in the brain. Decades of research using such brain–machine interfaces have shown that animals (non-human primates and rodents) can be operantly conditioned to self-regulate neural activity in various motor-related structures of the brain. Here, we ask whether the human brain, a complex interconnected structure of over 80 billion neurons, can learn to control itself at the most elemental scale—a single neuron.

We used the unique opportunity to record single units in 11 individuals with epilepsy to explore whether the firing rate of a single (direct) neuron in limbic and other memory-related brain structures can be brought under volitional control. To do this, we developed a visual neurofeedback task in which participants were trained to move a block on a screen by modulating the activity of an arbitrarily selected neuron from their brain.

Remarkably, participants were able to volitionally modulate the firing rate of the direct neuron in these previously uninvestigated structures. We found that a subset of participants (learners), were able to improve their performance within a single training session. Successful learning was characterized by (i) highly specific modulation of the direct neuron (demonstrated by significantly increased firing rates and burst frequency); (ii) a simultaneous decorrelation of the activity of the direct neuron from the neighbouring neurons; and (iii) robust phase-locking of the direct neuron to local alpha/beta-frequency oscillations, which may provide some insights in to the potential neural mechanisms that facilitate this type of learning.

Volitional control of neuronal activity in mnemonic structures may provide new ways of probing the function and plasticity of human memory without exogenous stimulation. Furthermore, self-regulation of neural activity in these brain regions may provide an avenue for the development of novel neuroprosthetics for the treatment of neurological conditions that are commonly associated with pathological activity in these brain structures, such as medically refractory epilepsy.

1 Krembil Brain Institute, Toronto Western Hospital (TWH), Toronto, Ontario M5T 1M8, Canada
2 Institute of Biomedical Engineering, University of Toronto, Toronto, Ontario M5S 3G9, Canada
3 Faculty of Medicine, University of Toronto, Toronto, Ontario M5S 1A8, Canada
4 Center for Advancing Neurotechnological Innovation to Application (CRANIA), Toronto, Ontario, M5G 2A2, Canada
5 The KITE Research Institute, University Health Network, Toronto, Ontario M5G 2A2, Canada
6 Electrical and Computer Engineering, University of Toronto, Toronto, Ontario M5S 3G4, Canada
7 Division of Neurosurgery, Department of Surgery, University of Toronto, Toronto, Ontario M5S 1A1, Canada
8 Max Planck-University of Toronto Center for Neural Science and Technology, Toronto, Ontario M5S 3G9, Canada

Received April 01, 2021. Revised August 16, 2021. Accepted September 03, 2021. Advance access publication October 8, 2021

Correspondence to: Kramay Patel
10 Navy Wharf Court, Toronto, ON M5V3V2, Canada
E-mail: kramay.patel@mail.utoronto.ca

Keywords: instrumental conditioning; neurofeedback; brain–machine interfaces; limbic; epilepsy

Abbreviations: BMI = brain–machine interface; LFP = local field potential; MD = modulation depth; RMS = root mean squared; SFC = spike field coherence

Introduction

Advances in physical and computational tools continue to inspire the development of devices to interrogate brain circuits and restore lost neural functioning. While the motor system has long been a target for such devices, there is an emerging interest in neuromodulatory as well as neuroprosthetic technologies for the interrogation and augmentation of cognition—in particular memory.[1–5] The seminal work of Eberhard Fetz in the late 1960s demonstrated that with the appropriate feedback and reward, monkeys can learn to control the activity of individual neurons in the primary motor cortex.[6,7] More recent work using advanced imaging and stimulation technologies in transgenic mice has demonstrated intentional neuroprosthetic learning of individual neurons within primary motor and visual cortices.[8–12] Whether such high-fidelity neuroprosthetic skill learning can be acquired in memory-related structures of the human brain remains unknown.[13,14]

At large spatial scales, scalp EEG has provided varied, albeit supportive literature regarding the efficacy of biofeedback to control oscillatory power in non-motor regions of the human brain.[15–19] On a mesoscopic scale, intracranial EEG recordings have shown humans can control oscillations in the local field potential (LFP) within medial temporal lobe structures.[20,21] Few have even reported the possibility of controlling neuronal activity in human medial temporal lobe,[22] and other motor-related structures[23]; however, such control relied on invocation of previously identified concepts or motor imagery. Thus, it remains unknown whether operant conditioning of individual neurons within memory structures of the human brain is possible.

To explore this question, we exploited the unique opportunity to obtain human single neuron recordings from limbic and other memory-related brain structures of epilepsy patients undergoing diagnostic depth-electrode recordings. We developed a closed-loop real-time instrumental learning task, where visual feedback is provided to participants as they try to learn to increase the firing rate of an arbitrarily chosen neuron. We show that: (i) humans can volitionally increase the firing rate of arbitrary individual neurons in these regions; (ii) as with all forms of instrumental learning, only a subset of participants learn the task (learners); and (iii) only learners demonstrated an increase in local spike field coherence (SFC), with the strongest SFC in the beta band (an oscillation not commonly investigated in these regions). Our findings show that: (i) instrumental learning to control arbitrary individual neurons is possible in the human brain; (ii) that such learning is possible outside of primary motor and sensory areas, and of particular interest, in mnemonic structures; and (iii) intriguingly, the unique beta band SFC signature of learners may provide insights into the neural mechanisms that facilitate this type of instrumental learning within these mnemonic structures of the human brain.[14,24]

Materials and methods

Electrode localization and electrophysiology

Participants were implanted with commercially available depth electrodes (Behnke–Fried Macro Micro, ADTech Inc) that have 8–10 macroelectrodes along the electrode shaft, and a bundle of eight microwires splayed out from the tip, plus one ground/reference microwire. The number and location of the electrodes varied between participants and were determined solely based on the clinical hypothesis of the epileptogenic zones(s). Electrode localization was performed by coregistering preoperative MRI with postoperative CT using the iELVIS toolbox.[25] Following localization, the precise location of each of the electrodes was determined, and later verified by a neurosurgeon.

Electrodes were connected to the Neuralynx Atlas Data Acquisition System (Neuralynx Inc, Bozeman, MT). The macroelectrodes were sampled at 4 kHz with a 16-bit resolution and bandpass filtered in hardware between 0.1 and 1 kHz. The microelectrodes were sampled at 32 kHz with a 16-bit resolution and were bandpass filtered in hardware between 0.1 and 8 kHz. A four-contact subgaleal electrode was used for ground and reference and was placed over the parietal midline facing away from the brain. The microelectrodes were referenced locally, to one of the eight wires within the same bundle. Neural data were synchronized with behavioural data using TTL triggers sent over Neuralynx's NETCOM protocol.

Online and offline spike detection and sorting

Online spike detection/sorting was used to drive the neurofeedback task (see the 'Neurofeedback task' section). Microelectrode channels were bandpass filtered between 300 and 3000 Hz,[26] and thresholded at five times the root mean squared amplitude. Channels in non-motor regions with well-isolated spikes were sorted using the KlustaKwik toolbox in SpikeSort3D (Neuralynx Inc., Bozeman, MT). Sorted templates for each channel were then sent back to the recording software, Pegasus (Neuralynx Inc., Bozeman, MT). The instantaneous timestamps of sorted neurons were then streamed to custom-written scripts in MATLAB (Mathworks Inc., Natick, MA) over the Neuralynx NETCOM protocol. Although online spike sorting was performed on numerous channels, only the online spike train of the direct neuron (i.e. the neuron being trained) was used for *post hoc* analysis. The remainder of the neurons (i.e. the indirect neurons) were detected and sorted offline.

For offline spike detection, all microelectrode channels were bandpass filtered between 300 and 3000 Hz, and spikes were subsequently detected using threshold crossings of a local energy measurement, calculated by convolving the raw signal with a kernel of approximate width of an action potential. All detected spikes were sorted using the semiautomatic template-matching algorithm OSort, which is available as open source.[27] Similar to previous work,[28] we classified clusters as putative single neurons

by looking at the following criteria: (i) violation of refractory period; (ii) shape of the interspike interval distribution; (iii) shape of the waveform; and (iv) separation from other clusters. Clusters that appeared similar to one another were merged, and clusters that were either contaminated with noise or failed the meet the criteria described were rejected. For the accepted clusters, the individual waveforms, along with the timestamp of each spike and its cluster definition were saved.

Neurofeedback task

We developed a novel, low-latency intracranial neurofeedback task in which the vertical movement of a red square on a screen was controlled by the smoothed instantaneous firing rate of a well-isolated neuron (called the direct neuron) from a mnemonic, non-motor structure in the human brain (Table 1). The direct neuron was chosen by the experimenter such that it was (i) well isolated from background noise; (ii) not contaminated by movement artefacts; (iii) had $<3\%$ of interspike intervals below 3 ms[27]; and (iv) had a baseline firing rate >0.5 Hz. The spike timings of the chosen, online sorted direct neuron were streamed to custom scripts in MATLAB using the Neuralynx NETCOM protocol. The spike timings were used to create a spike train, which was stored in a 2-s first-in-first-out buffer that was updated every ~40 ms. This spike train was then smoothed by convolving it with a 200-ms Gaussian kernel. The instantaneous value of the smoothed spike train was used as the control signal. Our visual neurofeedback task was developed using Psychtoolbox, and consisted primarily of a red square capable of moving vertically on a screen, along with a white horizontal line indicating the target (Fig. 1). The vertical movement of the red square was controlled by the described control signal. Each participant was asked to modulate their brain activity to move the square above the target line and hold it there for at least 0.5 s. Every time the block went above the target line, its colour was changed to purple to clearly identify to the participant that the block was above the line.

Each testing session began with a ~4-min baseline session in which participants stared at a fixation cross for 1 min, followed by a 1-min eyes-closed period, followed by passively viewing a red square moving on the screen for 2 min. All participants were instructed to minimize body movements throughout the session, and any trials with overt physical movements were noted and later rejected. Following the baseline session, an appropriate well-isolated neuron with a firing rate above 0.5 Hz was chosen as the neuron to be trained. If the individual participated in more than one training session, we ensured to choose a neuron from a different electrode, to ensure independence of each session. Using this well-isolated neuron, we then performed a short, ~5-min training and familiarization session in which we asked the participant to modulate their brain activity to move the block on the screen. This session was also used to determine an appropriate starting difficulty for the task (i.e. the position of the horizontal target bar on the screen). An appropriate difficulty level was defined as a multiple of the standard deviation of the baseline firing rate of the selected neuron, chosen such that the participants required ~60 s to achieve success. To prevent biasing each participant's unique control strategy, we never provided the participants with explicit instructions on how to modulate their neural activity. Furthermore, literature suggests that participants who report no specific control strategies demonstrate better neurofeedback performance than those who report specific mental strategies.[18,29]

Following the familiarization/training portion, we moved onto the primary testing portion of the session. This portion was divided into several blocks, consisting of 10 trials each. Each trial began with a start message, following by a fixation cross (3 s), followed by the actual trial in which the participants volitionally moved the red square on the screen. Each trial continued until the participant was successfully able to push the block above the line and hold it there for 0.5 s. This resulted in a 'Trial Successful' message (1 s) followed by a distractor math question (i.e. addition of three integers between 1 and 5). Answering the question triggered the start of the next trial. Participants were instructed to try and complete each block of trials (i.e. 10 trials) in 10 min or less. If the participants successfully completed all 10 trials in <10 min, the difficulty of the next trial was marginally increased (by increasing the target threshold by 0.2 standard deviations). Participants were required to complete at least three blocks of testing, or 30 trials. Sessions in which testing was interrupted before the minimum number of trials were completed were not used in the analysis. Testing was continued up to a maximum of 12 blocks, or stopped earlier if otherwise interrupted by visitors, clinical interventions or self-reported fatigue. After the testing session, a post-testing baseline session was performed, which mirrored the pretesting baseline session described previously.

Table 1 Participant demographics

Session ID	Patient ID	Sex	Handedness	Learner session?	Direct neuron anatomy	Hemisphere	Trials completed, *n*	Age
1	1	Male	Right	Yes	Lateral orbitofrontal	Left	50	25.1
2[a]	1	Male	Right	Yes	Lateral orbitofrontal	Left	53	25.1
3	2	Female	Right	No	Amygdala	Left	39	46.1
4	3	Female	Right	Yes	Hippocampus	Right	117	20
5	3	Female	Right	No	Insula	Right	77	20
6	4	Female	Left	Yes	Hippocampus	Left	77	25.5
7	5	Male	Right	Yes	Amygdala	Right	30	26.2
8	5	Male	Right	Yes	Amygdala	Left	39	26.2
9	6	Male	Right	Yes	Lingual cortex	Left	59	38.1
10	6	Male	Right	Yes	Hippocampus	Left	59	38.1
11	7	Male	Right	Yes	Hippocampus	Left	94	28.5
12[a]	8	Female	Right	No	Hippocampus	Left	48	25.9
13	9	Male	Right	No	Hippocampus	Right	48	24.6
14	9	Male	Right	No	Amygdala	Right	52	24.6
15	10	Male	Right/Left	No	Hippocampus	Right	29	25.3
16	10	Male	Right/Left	Yes	Hippocampus	Right	48	25.3
17	11	Male	Right	No	Amygdala	Right	47	28.1

[a]Missing macro data.

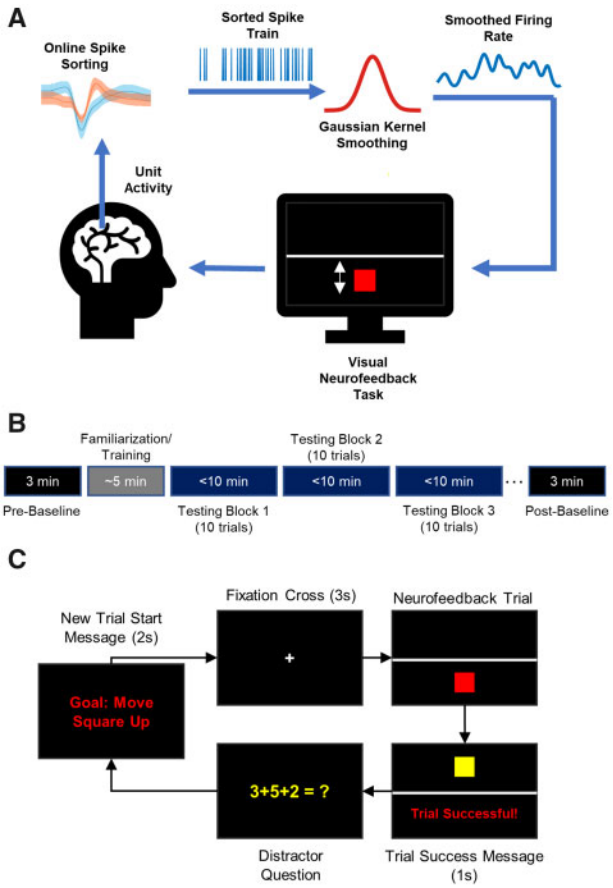

Figure 1 Visual neurofeedback task for modulating single neuron activity in the human brain. (**A**) Schematic showing the overall setup of the neurofeedback task. Single unit activity is extracted from the microwires in implanted Behnke–Fried Macro Micro electrodes (AdTech) using the Neuralynx Atlas Digital Lynx system. Neurons from relevant channels were sorted online (using templates created with the KlustaKwik algorithm), and streamed over the Neuralynx NETCOM protocol to custom scripts in MATLAB. These streamed, online spike trains were smoothed using a 200-ms Gaussian kernel to extract the instantaneous firing rate, which was then used as the control signal for the task. The neurofeedback task itself involves a red square moving vertically in response to the instantaneous firing rate of the chosen neuron. A horizontal white line indicates the target threshold, the crossing of which for 0.5 s results in a successful trial. (**B**) Schematic showing the experimental design, and order of a single testing session. An initial baseline session was performed, followed by 5 min of familiarization/training (during which an appropriate task difficulty is chosen). Following this, the testing phase began, which consists of at least three blocks (of 10 trials each). Each trial ended with a success, and each block was required to be finished in 10 min or less. Testing continued until a maximum of 12 blocks. Following testing, another baseline session was performed. (**C**) Schematic showing a single testing trial. Each trial began with a start message indicating the goal, followed by a fixation cross, followed by the actual trial. Each trial ended with a success message, followed by a quick distractor question involving the addition of three small integers.

Data analysis

All analyses were performed in MATLAB with custom-written routines and with the aid of open-source toolboxes where possible. All LFP data were low-pass filtered at 250 Hz, downsampled to 1 kHz and notch filtered at 60, 120, 180 and 240 Hz. Online and offline spike times were used to create binary spike trains with a bin width of 1 ms (corresponding to 1 kHz sampling). Signal to noise ratio (SNR) was calculated as the peak-to-trough amplitude of the average neuron waveform, divided by the standard deviation of the waveform residuals after the mean waveform was subtracted.[30]

Firing rates were calculated by convolving the binary spike trains with a scaled 200 ms Gaussian ($\sigma = 200$ ms). Average firing rates correspond to the mean of the smoothed spike train in each trial. Peak firing rate corresponds to the maximum of the Gaussian convolved smoothed spike train in each trial. Modulation depth (MD) of each neuron was calculated as the average firing rate in the 1-s window before success minus the average firing rate in the 1-s window after success. Bursts were detected using a non-parametric version of the commonly used Poisson-surprise method, called the rank-surprise method,[31] which identifies bursts on the basis of the probability that a given number of spikes occur in a given duration, given the overall distribution of the interspike intervals for a particular neuron. Burst frequency was defined as the total number of bursts in each trial divided by the duration of each trial. A session was marked as 'Learner' if a linear regression of the average or peak firing rate with the trial number resulted in a significant positive slope. i.e. 'Learner' sessions were those in which there was a significantly positive trend in the average and/or peak firing rates within a single testing session. All other sessions were marked as 'Non-Learner'. In all analyses, the term 'early' corresponds to the data from the first 15 trials of the testing session, and 'late' corresponds to data from the last 15 trials of the testing session. This definition was chosen since the minimum number of trials completed by each participant was 30. Correlation of the direct neurons with the neighbouring indirect neurons was examined by calculating Pearson's correlation coefficient between the binary spike train of the direct neuron, and the binary spike trains of all neighbouring indirect neurons. For each trial, the pairwise correlations between the direct neuron and all neighbouring indirect neurons were averaged. Early to late changes in correlation (i.e. $\Delta_{correlation}$) were calculated by taking the average pairwise correlation coefficient in the 'late' trials and subtracting the average pairwise correlation coefficient in the 'early' trials.

All spectral analyses were performed using the open-source Chronux toolbox (http://chronux.org) for MATLAB. Spectral estimates were obtained using a multi-taper method, with a total of five tapers and a time-bandwidth product of three. For all spectrograms, a moving window size of 1 s was used (to accurately resolve frequencies as low as 1 Hz), with a step size of 50 ms. SFC was calculated as shown in Equation (1):

$$C = \frac{R_{xy}}{\sqrt{R_{xx}}\sqrt{R_{yy}}} \quad (1)$$

where R_{xx} and R_{yy} are the power spectra of the spike train and LFP oscillation, respectively, and R_{xy} is the cross-spectrum. All local SFC values were calculated between any given putative neuron and the closest macroelectrode to ensure that the spiking activity itself did not contaminate the power spectrum of the recorded oscillations. Since SFC estimates can be affected by the firing rate of the selected neuron, we performed a probabilistic spike thinning procedure to equate the firing rate of the early and late trials.[32] To do so, we convolved the binary spike trains with a 10-ms Gaussian kernel, and averaged them across trials. Next, we determined the probability that a spike should be removed at any given point in time by subtracting the early and late firing rates at each point in time, and dividing by the maximum rate in the early and late trials at that point in time. Using this probability train, we randomly removed spikes from the original spike trains. For example, if the probability value at any given point in time was 50%, then the probability that a spike in the late trials at that point in time was removed was 50%, resulting in roughly 50% of the spikes at that

point in time being removed. We verified that the spike thinning procedure successfully eliminated any differences between the firing rate between the early and late trials (Supplementary Fig. 6).

Statistical analysis

Unless otherwise specified, the bold line in all figures corresponds to the mean of the data, and the shaded area and/or the error bars correspond to the SEM. In all box plots, the shaded box corresponds to the SEM, and the bold vertical line corresponds to the standard deviation. All means, standard errors and statistical tests were calculated across sessions. The number of participants or sessions contributing to each figure is indicated with the corresponding *n*-value. Wherever necessary, data outliers were removed using the Grubbs outlier test. Wherever possible (given data normality and absence of outliers), parametric tests were used to test for significance. Otherwise, non-parametric equivalents were used. When determining significant fluctuations in averaged waveforms (time-aligned or power spectra), a non-parametric permutation test was using with random time shuffling and 2000 iterations. The specific statistical test used for each figure is stated clearly in the text and/or in the figure legend. The data epoch being analysed in any particular figure is clearly labelled and described in the figure legend. Significance for all statistical tests was set at $P < 0.05$. A single asterisk indicates significance at the $P < 0.05$ level, double asterisk indicates significance at the $P < 0.01$ level and a triple asterisk at the $P < 0.001$ level. All statistical tests were performed using either custom scripts or built-in functions in MATLAB.

Data availability

The spike detection and sorting toolbox used for offline sorting (OSort), and the Chronux toolbox (use for spectral analysis) are both available as an open-source toolboxes. Data and custom MATLAB scripts used for analysis here are available on reasonable request from the corresponding author.

Results

Performing the neurofeedback task

We developed a neurofeedback task that required upregulation of the firing rate of an arbitrarily chosen neuron (henceforth called the direct neuron) (Fig. 1 and Table 1; see the 'Materials and methods' section for details on the choice of the direct neuron). Direct neurons were well isolated (SNR = 4.76 1.38) and largely free of external noise (% of the interspike intervals below 1 ms = 0.09 0.19%). Spiking activity of the direct neuron was detected and sorted in real time. Online spike trains were convolved with a 200 ms Gaussian[21] to obtain its smoothed instantaneous firing rate. The smoothed firing rate of the direct neuron was linearly mapped onto the vertical position of a square on a screen placed in front of the participant (see 'Materials and methods' section). Participants were instructed to try and move the block above a white horizontal line (threshold). Maintaining the box above threshold for over half a second indicated success. A successful trial message was displayed, followed by a distractor question, after which the next trial was triggered. In this way, it was ensured that each trial ended in a success. Testing was divided into blocks of 10 trials, and the participants were asked to finish the 10 trials in 10 min or less. To keep the participants motivated, we increased the difficulty of the next block of trials (by moving up the target line) if the previous 10 trials were completed in <5 min. Learning was evaluated *post hoc* by analysing changes in the average and peak firing rate of the trained neuron as a function of the trials. Eleven participants completed a total of 17 sessions, where they controlled a different direct neuron in each session. Since each session involved controlling a different neuron, each session was treated as an independent session.

All participants completed at least 30 trials (57 22 trials) in which they upregulated the activity of the direct neuron to complete each trial (Fig. 2A and C, top). Interestingly, the firing rate of the population of other neurons recorded from the same bundle of microwires as the direct neuron (henceforth called indirect neurons) did not change before successful completion of the trial (Fig. 2A). Similarly, we did not observe a significant change in the firing rate of the 743 indirect neurons recorded from other microwire bundles throughout the brain before successful trial completion (Supplementary Fig. 1A). The indirect neurons were also well isolated (SNR = 3.98 1.72) and largely free of noise (%interspike intervals below 1 ms = 0%). To further quantify task-contingent changes in firing rate, we calculated the MD of the direct neurons,[10] defined as the average firing rate in the one second window before success minus the average firing rate in the 1-s window after success. If success were triggered by random bursts of activity, the MD would be close to zero since the bursts would probably continue into the post-success period. To the contrary, we saw a sharp decline in the activity of the direct neurons immediately following success (Fig. 2B), resulting in a MD significantly >0 ($P < 0.001$, single sample *t*-test). To determine whether this type of upregulation was specific to the direct neuron, we calculated the MD of indirect neurons. Indirect neurons' firing rates were neither task contingent, nor modulated with the direct neuron as evidenced by their MD being close to zero (Fig. 2B and Supplementary Fig. 1). While the population of indirect neurons as a whole was not modulated by the task, it may be possible that individual indirect neurons were task contingent (i.e. comodulated along with the direct neuron). To determine whether an individual indirect neuron was significantly modulated by the task, we calculated the average MD of each indirect neuron in a 2-s window centred on trial completion (same as before for direct neurons). We then generated a null distribution of average MDs for each neuron by randomly sampling time points matching the number of trials in each session, calculating the MD in a 2-s window around each time point and averaging these MDs. This was repeated 1000 times to generate a null distribution. An indirect neuron was labelled as significantly modulated if its actual MD was >97.5th or <2.5th percentile of the null distribution of MDs. In doing so, we found that while most indirect neurons (101/116 or 87.1%) did not show any significant modulation around trial completion, a small fraction (15/116, or 12.9%) had firing rates that were significantly modulated around trial completion. When restricting this analysis to the 'late' trials (i.e. the last 15 trials), an even lower number of neurons (10/116, or 8.6%) appeared to be modulated around trial completion. We also simulated the task using the activity of each indirect neuron to determine what percentage of indirect neurons could complete each trial before the direct neurons. Although a relatively large proportion of indirect neurons were upregulated before the direct neuron in the early stages of learning, successful learning decreased this proportion significantly, pointing to the acquired volitional and specific control over the activity of the direct neuron (Supplementary Fig. 3).

Learning to improve performance

To keep participants motivated we 'staircased' the difficulty of the neurofeedback task, i.e. increased the difficulty as the participants became more adept at the task. Thus, instead of using time-to-success as a primary measure of learning, we evaluated learning by observing changes in the firing rate of the direct neuron over the

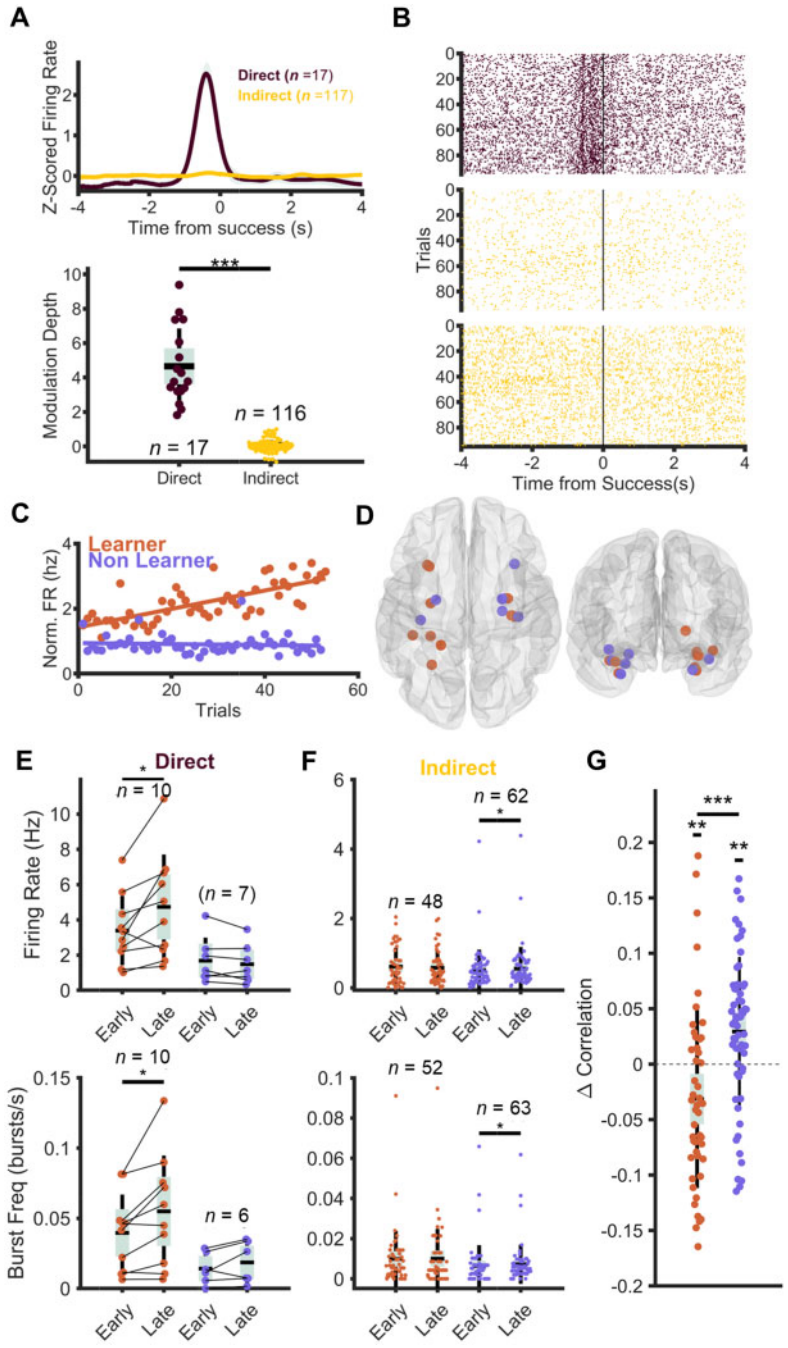

Figure 2 Learning to upregulate the activity of direct neurons using neurofeedback. (**A**) *Top*: Firing rate of the direct neuron increased sharply immediately before success ($t = 0$) peaking at 380 ms before success, and returned to baseline immediately after success. Firing rate of indirect neurons was not modulated in the same epoch around success. *Bottom*: Modulation depth of direct neurons is significantly >0 ($P < 0.001$, single sample two-tailed t-test) and significantly greater than that of indirect neurons ($P < 0.001$, independent samples t-test). Outliers are removed using the Grubbs method. (**B**) Representative spike rasters across all trials from a representative session. *Top* (black) panel shows the spike raster of the direct neuron, and *bottom* two (yellow) panels show the spike rasters of neighbouring indirect neurons recorded from the same bundle as the direct neuron. (**C**) Representative learner and non-learner sessions. A significant positive slope in the regression line between trial number and peak or average firing rate in each trial results was considered a learner session. All other sessions were defined as non-learner sessions. (**D**) Anatomical distribution of the direct neurons, colour-coded to match whether the session was a learner or non-learner. (**E**) Changes in the firing rate (*top*) and burst frequency (*bottom*) of direct neurons within a single session grouped by learner and non-learner sessions. Firing rate increased from the early to late trials in the learner sessions ($P = 0.024$, paired t-test, Cohen's $D = 0.8591$), and so does burst frequency ($P = 0.019$, paired t-test, Cohen's $D = 0.90$). Changes in the firing rate ($P = 0.21$, paired t-test, Cohen's $D = 0.54$) and burst frequency ($P = 0.20$, paired t-test, Cohen's $D = 0.61$) are not evident for the non-learner sessions. (**F**) Same as in **E**, but for indirect neurons. Firing rate and burst frequency do not change in learner sessions (firing rate: $P = 0.99$, Wilcoxon Signed-Rank Test, common language effect size = 0.52); burst frequency: $P = 0.98$, paired t-test, Cohen's $D = 0.0034$), but increase significantly for the non-learner sessions (firing rate: $P = 0.018$, paired t-test, Cohen's $D = 0.31$; burst frequency: $P = 0.045$, paired t-test, Cohen's $D = 0.26$). (**G**) Change (late minus early) in spike correlations between the direct neuron and the neighbouring indirect neurons recorded from the same bundle of microwires. Correlations with neighbouring neurons decreased significantly in the learner sessions ($P = 0.0094$, single sample two-tailed t-test), and increased significantly in the non-learner sessions ($P = 0.0013$, single sample two-tailed t-test). Change in correlation is significantly different between the learner and non-learner sessions ($P < 0.001$, independent samples t-test). $^{*}P < 0.05$, $^{**}P < 0.01$ and $^{***}P < 0.001$; n indicates the number of neurons.

course of each session.[20] 'Learner' sessions were defined as those sessions where there was significant upregulation of the average and/or peak firing rate of the direct neuron as a function of time during the session. To do this, we performed a linear regression between the average and peak firing rate of the direct neuron as a function of trial number. Sessions were defined as learner sessions if there was a significantly positive trend in either the peak or average firing rate of the direct neuron (Fig. 2D; see 'Materials and methods' section for more details). With this definition, we defined 10 sessions as learners (across seven patients) and the remaining seven sessions as non-learners (see Table 1 for participant demographics). It should be noted that, by definition, every trial ended in a success. However, the success could be triggered by a random, spontaneous burst of activity of the direct neuron (as is the case throughout the non-learner sessions) or a volitionally driven upregulation of the direct neuron firing (as is the case in the learner sessions). In the non-learner sessions, individuals are unable to volitionally upregulate the activity of the direct neuron, thus there is no increase in the average firing rate or burst frequency of the direct neuron (Fig. 2C, E and F). Early in the learner sessions, success is probably driven by random bursts of activity, but as learning occurs, individuals are able volitionally modulate the activity of the neuron to increase its firing rate (Fig. 2E) as well as burst frequency (Fig. 2F) to complete the trials more rapidly (Supplementary Fig. 2). These findings agree well with the general observation that neuroprosthetic skill learning is not uniformly acquired.[20,33,34] While the 'staircased' difficulty in this task made it impossible to analyse the time-to-completion over the entirety of each session, we investigated sessions that had at least 20 consecutive trials where the difficulty was kept the same. We analysed the time-to-completion in these contiguous trials using a non-linear, exponential decay regression model and found that the time-to-completion decayed in learner sessions, but not in non-learner sessions (Supplementary Fig. 2). Thus, while all participants were able to upregulate the activity of the direct neurons, only during some sessions were they able to improve their performance in the task. Although we did not have sufficient data to interrogate the role of anatomic specificity in successful and unsuccessful learning, we did investigate whether the hemispheric localization of the direct neuron was correlated with the ability to 'learn' the task. To do this, we performed a chi-square test for association between learning (yes versus no) and hemisphere (left versus right). There was no significant association between learning and hemispheric localization of the direct neuron (Pearson chi-square value: 2.837, $P = 0.092$, Cramer's $V = 0.408$). Furthermore, we also asked whether the cellular morphology and/or firing characteristics of the direct neuron influences the ability to learn the task. To do this, we performed a binary logistic regression with learning outcome (i.e. learner versus non-learner) as the response/dependent variable and various cellular morphology and firing characteristics as the covariates (see Supplementary Fig. 4 for details). Although this model was significant, the only meaningful predictor of learning was the baseline firing rate of the direct neuron ($B = 2.246$, odds ratio = 9.45, $P = 0.071$). Separately, we found that the baseline firing rate of the direct neuron in the learner sessions was significantly higher than the baseline firing rate in the non-learner sessions, providing further evidence that direct neurons with higher baseline firing rates may be easier to further upregulate (see Supplementary Fig. 4 for details).

As expected by the definition of the learner and non-learner groups, the average firing rate of the direct neuron in learner sessions was significantly higher in the later trials (i.e. the last 15 trials) compared to the early trials (i.e. the first 15 trials) (Fig. 2F, top). The burst frequency of the direct neurons (calculated using a modified Poisson-surprise method) also increased significantly in the learner sessions but not in the non-learner sessions (Fig. 2F, bottom). Indirect neurons demonstrated the opposite trend, with average firing rate and burst frequency increasing (by a small albeit significant magnitude) in the non-learner sessions, but not in the learner sessions (Fig. 2G). The firing rate or burst frequency of indirect neurons recorded from other brain regions did not change from early to late trials in learners or non-learners (Supplementary Fig. 1C and D). These data reveal a stark dissociation between neural activity of the direct and indirect neuronal populations, where during learner sessions, participants selectively modulate the activity of the direct neurons and in non-learner sessions they unknowingly modulate the activity of neighbouring neurons, while failing to specifically modulate the direct neuron. Thus, learning is accompanied by selective, volitional control over the direct neuron, whereas unsuccessful learning is characterized by non-specific modulation of the entire neural subpopulation consisting of both direct and indirect neurons. This dissociation is further exemplified by the decorrelation of the activity of the direct neuron from the neighbouring indirect neurons in the learner sessions, and an increase in this correlation in the non-learner sessions (Fig. 2H). This finding mirrors similar findings reported using calcium imaging studies in rodents.[8,9]

Spike field coherence develops during learning

During similar neuroprosthetic skill acquisition in rodents, learning is accompanied by increased corticostriatal communication evidenced by corticostriatal coherence observed in the LFP,[11] as well as SFC between cortical neurons and striatal oscillations and vice versa.[10–12] Striatum is not a clinical target in intracranial EEG recordings in epilepsy patients, and thus we were unable to test the hypothesis of striatal communication in the volitional control of individual neurons in humans. However, we used rodent SFC findings to motivate a similar analysis to infer 'communication through coherence'[35] if such SFC was observed. Towards this end we computed the SFC between direct neurons and LFPs recorded by the closest macro contact of the Behnke–Fried electrode (local LFP; Fig. 3A). For learners we found a significant increase in the SFC in the 10–20 Hz range immediately before success (Fig. 3A and B), while the non-learner population displayed no such increase in SFC. The differences were most pronounced in the 12–15 Hz frequency range (Supplementary Fig. 5). The ability to learn this skill is thus associated with a unique electrophysiological state of the brain,[36] evidenced by increased SFC in the beta-frequency range that is different from other 'learning' states of the human brain where theta appears to play an almost exclusive role.[37–40]

Since most indirect neurons are not task relevant (their firing rates did not contribute to success), we anticipated that these indirect neurons would not develop the same learning-related SFC that was observed for direct neurons. To test this hypothesis, we calculated the SFC for indirect neurons to the local LFP and found no learning-related changes in both the learner and non-learner populations (Fig. 3C). Before calculating the SFC, the firing rate of early and late trials were matched using a spike thinning procedure to prevent any biases resulting from unequal firing rates between the conditions (Supplementary Fig. 6, see the 'Materials and methods' section for details). Phase-related measures can often be affected by changes in oscillatory power.[41] To determine whether the observed change in learning-related SFC was affected by spectral power changes, we computed the power spectra in the same time period in early and late trials. We found no differences in the power in the 10–20 Hz frequency bands (Fig. 3D and E). We have shown that visually evoked responses are readily seen in the medial temporal lobe and associated structures[42] that may confound SFC estimates. To address this, we subtracted the event-

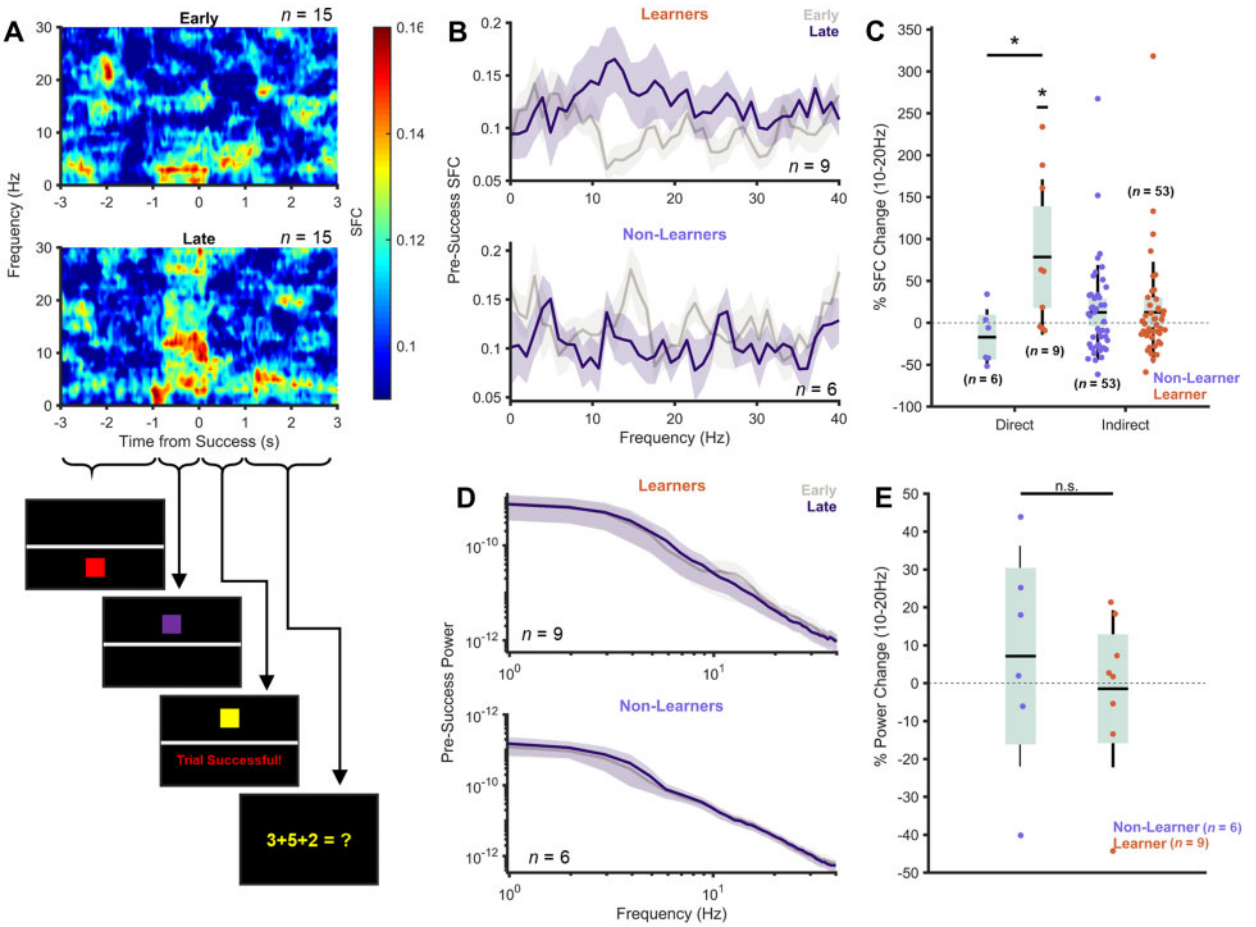

Figure 3 Local SFC in the 10–20 Hz range emerges as learning progresses. (**A**) Coherogram of grand-average SFC of the direct neurons in early (*top*) and late (*bottom*) trials. Schematics of the visual state of the task during each of the relevant periods (*below*). Note the significant increase in the SFC in the 10–20 Hz band immediately preceding success (see the Supplementary material for details regarding the SFC calculation). (**B**) The grand-average SFC in a 1-s window immediately preceding success. Notice the substantial increase in the SFC in the 10–20 Hz range in the learner sessions, but not in the non-learner sessions. (**C**) Percentage (%) change in SFC in the 10–20 Hz range (from early to late trials) in the learner and non-learner sessions, for direct and indirect neurons. For the direct neurons, SFC increased significantly in the learner sessions ($P = 0.035$, single sample two-tailed *t*-test) but not in the non-learner sessions ($P = 0.26$, single sample two-tailed *t*-test). The % change in SFC was also significantly higher in learner sessions compared to non-learner sessions ($P = 0.032$, independent samples *t*-test). For the indirect neurons, there is no change in the learner ($P = 0.76$, Wilcoxon Sign Rank Test) or the non-learner sessions ($P = 0.37$, Wilcoxon Sign Rank Test). (**D**) Grand-average power spectra in the 1-s window immediately preceding success for learner (*top*) and non-learner (*bottom*) sessions. There were no significant learning-related changes in the power spectrum in the presuccess interval, in the 10–20 Hz frequency range. (**E**) Percentage power change in learner and non-learners is not different from 0 ($P = 0.85$ learners, $P = 0.57$ non-learners, single sample two-tailed *t*-test), or from each other ($P = 0.53$, independent samples *t*-test). *$P < 0.05$.

related potential for each LFP channel from each trial before calculating the coherence values. In doing so, the early to late differences in the 10–20 Hz SFC persisted in the learner sessions (Supplementary Fig. 7), providing further evidence that the SFC observed here was not artefactual in nature. Thus, in the absence of firing rate changes, power-related changes and evoked responses, the observed changes in the learning-related SFC are related to increased precision of spike timing immediately before success.

Additionally, we observed instances where the same participant could learn successfully in one session, but not in another (Table 1). Despite this, we observed the learning-related SFC changes confined to the learner sessions, suggesting the specificity of these changes to the act of learning itself, and not to other demographic factors.

The observed SFC in the beta band might be due to volume conducted low frequency oscillations. To address this specifically we calculated the SFC between the direct neurons and the LFP at non-local macroelectrodes throughout the brain (Supplementary Fig. 8). In addition to the increase in the 10–20 Hz frequency band SFC between the direct and non-local LFP (Supplementary Fig. 8D), there was an even more profound increase in theta frequency SFC between the direct neuron and the non-local LFP. Since theta is a ubiquitous oscillation in the human brain,[43] including the human hippocampus[37,42] and more likely to contribute to volume conduction,[44] the beta-frequency learning related SFC appears to be both specific in frequency range and local to the direct neuron during neuroprosthetic skill learning.

Learning related spike field coherence is distinct from anticipatory reward

Since the participants were asked to hold the square above the threshold for >500 ms, we wondered whether the observed SFC in the period immediately before success was driven by a reward anticipation mechanism.[45] To test this theory, we extracted epochs around unsuccessful threshold crossings, i.e. points in time when the firing rate of the direct neuron crossed the threshold but for an

insufficient time to trigger a successful trial. We hypothesized that if the 10–20 Hz SFC we observed in the presuccess period was indeed the result of an anticipatory reward mechanism, we would observe a similar increase in the 10–20 Hz SFC immediately after unsuccessful threshold crossings. Arguing against such an anticipatory reward mechanism, the spike-field coherogram of the threshold crossing-aligned epochs (Fig. 4A) did not demonstrate an increase in the 10–20 Hz SFC immediately after threshold crossings (Fig. 4B, top). Furthermore, there was no change in the SFC in this frequency band in the post-threshold crossing window between the early and late trials (Fig. 4C) for learners and non-learners, confirming that this type of reward anticipation does not drive the learning-related SFC changes observed in the success aligned epochs.

We did observe a significant increase in the delta-band (1–3 Hz) SFC immediately following threshold crossings (Fig. 4B). This finding was concordant with the increased delta SFC observed in the window immediately surrounding success (Fig. 3A). To determine whether this delta SFC was learning related, we compared the SFC in the post-threshold crossing window in the early versus the late trials (Fig. 4C). We observed no significant difference in the delta SFC in this window in the early versus late trials (for learners and non-learners), suggesting that the delta SFC was not learning related and was probably related to the design of the neurofeedback task. Consistent with this hypothesis the delta-SFC increase was associated with a delta power increase in a similar time window (Supplementary Fig. 9). This suggests that the observed delta-SFC increase is probably driven by the image onset evoked response due to the colour change of the square from red to purple when it crosses the threshold[42] rather than either anticipatory reward or a learning-related mechanism.

Discussion

Here we demonstrate, using a visual neurofeedback task, that humans can learn to upregulate the activity of arbitrarily chosen neurons in memory-related structures of their brain in a highly specific and volitional manner. Our results greatly extend non-human primate and rodent single neuron neuroprosthetic skill learning that has been explored in sensorimotor and associational structures, by showing uniquely that instrumental learning at the singe neuron level can occur in mnemonic structures. Additionally, these findings bridge an important gap in rodent and non-human primate neuroprosthetic skill learning and human learning, by showing that in humans such learning can occur in limbic and other mnemonic structures.

A large body of existing literature provides evidence for this type of neuroprosthetic skill learning in the motor cortex of rodents[9–12,46] and primates.[6,7,47–49] Early work by Fetz and colleagues showed that macaques can be operantly conditioned to control firing rates of individual units in the motor cortex.[6,7] Although this control was volitional, it was not highly specific to the trained neuron. They found that in about half of the recorded pairs of neurons (in which one neuron was trained and the other was simply recorded), activity was highly correlated during the task.[6] Contrary to this, we found that the volitional control of direct neurons in our task was highly specific (Fig. 2A and B). While we found a small fraction of recorded indirect neurons (12.9%) that were significantly modulated around trial completion, as a population, these neurons did not show significant modulation. These findings agree well with recent evidence from rodent *in vivo* calcium imaging studies.[8,50] Clancy and colleagues[50] used two-photon imaging in the motor and somatosensory cortex to show that as learning progresses, direct neurons develop coordinated activity patterns. While indirect neurons in the immediate vicinity of the direct neurons were modulated around target hits at the beginning

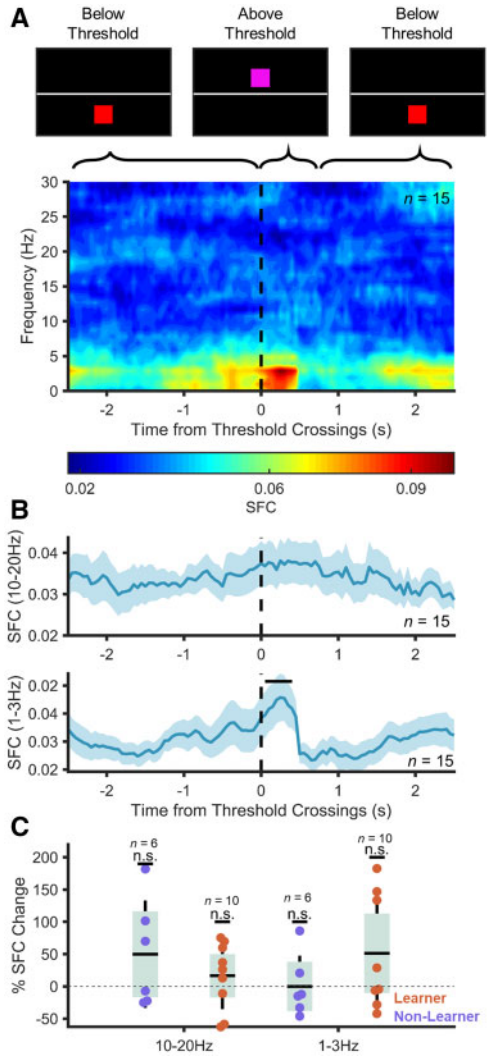

Figure 4 Unrewarded threshold crossings do not result in a change in SFC in the 10–20 Hz band. (**A**) Spectrogram showing the SFC between the direct neurons and the local LFP aligned to unrewarded threshold crossings. Notice that there is no significant increase in SFC in the 10–20 Hz range following the threshold crossings. (**B**) Average of SFC in the 10–20 Hz band (*top*) and 1–3 Hz band (*bottom*) across session (average traces were obtained by averaging the mean, band-limited SFC within each session, and then averaging these mean traces across sessions). Note that SFC does not increase in the 10–20 Hz band following threshold crossings, but it does increase significantly following threshold crossings in the 1–3 Hz band (significant portions indicated with a bold line on top of the graph, $P < 0.05$ non-parametric permutation testing with random time shuffles, 2000 iterations). (**C**) Change in SFC between early and late trials in the 10–20 Hz and 1–3 Hz frequency ranges for learners and non-learners. No significant changes observed (10–20 Hz: learner, $P = 0.37$, non-learner $P = 0.20$ single sample two-tailed *t*-test; 1–3 Hz: learner, $P = 0.45$; non-learner, $P = 0.99$, single sample two-tailed *t*-test). No learning-related changes in SFC observed in either frequency band following threshold crossing.

of training, this task-related modulation disappears as learning progresses. Prsa and colleagues,[8] who also used a similar behavioural model, corroborated these findings. In our work, similarly, we show that direct neurons become decorrelated with the local indirect neuronal population as learning progresses (Fig. 2H). Fascinatingly, this decorrelation only occurred in sessions where learning was successful. Our findings, along with the emerging animal literature, point to the ability of upstream circuits to 'hone in' on individual direct neurons and specifically modulate their

activity to efficiently modulate the feedback.[8,50,51] In conjunction with previous single-cell stimulation experiments that point to the ability of downstream circuits to read-out activity of individual neurons,[52] our results emphasize the importance and relevance of the activity of individual neurons.

Our findings, that only some neurons can be volitionally controlled (learner sessions), also agree well with the rest of the brain–machine interface (BMI) field.[20,34,36,53] Sadtler and colleagues[53] used a brain–computer interface in Rhesus macaques to beautifully demonstrate that animals could easily and robustly generate arbitrary activity patterns, as long as they were within the 'intrinsic manifold' of possible activity patterns. This intrinsic manifold comprises the characteristic, low-dimensional activity patterns of the local subpopulation of neurons, and is probably constrained by the inherent connectivity within this subpopulation. Generating activity patterns outside this manifold is more difficult and may require greater amounts of training.[36,53] Athalye and colleagues[34] also recently demonstrated that certain groups of mice can robustly modulate specific patterns of neural activity, whereas other mice cannot. This may also be attributed to the targeted activity falling outside of the intrinsic manifold for the non-learning mice. It is important to note that learning is also constrained by the duration of our sessions (which typically lasted <1 h). Non-learner sessions may turn into learner sessions given sufficient time and training. Presumably, the targeted activity of the direct neurons falls outside of the intrinsic manifold, and thus would require more time and training. In the context of motor control and motor skill learning, the intrinsic manifold can be thought of as the set of existing muscle synergies, which correspond to coordinated groups of muscles with specific amplitude weights.[54] Thus, learning a new motor skill is easier if the skill can be produced by a new combination of existing muscle synergies (i.e. learning within the intrinsic manifold). However, learning a motor skill that is incompatible with the existing set of muscle synergies is more difficult and may require more time (i.e. learning outside the intrinsic manifold).

In our work, we found select cases where the same subject could modulate activity of one neuron in a particular brain region, but not of another from the same region. This suggests that brain localization and patient-specific demographics (such as age, sex, years of education, number of anti-epileptic drugs, years of seizures, etc.) probably do not play a critical role in determining whether a session will be a learner or a non-learner. Instead, the ability to learn probably depends on the location of a neuron within an intrinsic manifold (as discussed previously), and also additional situational factors such as attention, motivation,[55] sleep and fatigue,[56] and mood.[57,58] We also found that learning may be influenced by the baseline firing rate of the direct neuron (Supplementary Fig. 4). Learner sessions typically had higher firing rates than non-learner sessions. A neuron with a high baseline firing rate may have a higher firing rate ceiling, making it potentially easier to upregulate. In other words, high firing rates of these neurons probably fall into the intrinsic manifold. Similar to this finding, Best and colleagues[59] previously reported that narrow spiking neurons with high firing rates consistently decoded motor parameters better than broad spiking neurons with lower firing rates.

In animals, control at the single neuron level in the motor cortex has been shown to require the dorsal striatum,[11,12] which serves as an input tier for the basal ganglia. Hence, neuroprosthetic skill learning in the motor cortex is largely analogous to motor learning, in which cortico-basal ganglia loops facilitate an action selection process where competing motor programs are either inhibited or released from inhibition. This is facilitated by parallel direct and indirect pathways that allow disinhibition and inhibition, respectively, of neuronal ensembles in the sensorimotor cortices, allowing for selection of a contextually relevant motor program.[60] Similar cortico-basal ganglia loops are implicated in selection and generation of a variety of different cognitive patterns that may facilitate more abstract skill learning.[61] In fact, recent rodent studies provide convincing evidence that animals can modulate highly specific neuronal activity in primary sensory cortex that again is dependent on the dorsal striatum, similar to learning in the motor cortex.[10] Since most of the neocortex projects to the dorsal striatum,[13,62] we anticipate that this type of neuroprosthetic skill learning may be possible in most of the neocortex.

In this study, however, we demonstrate that this type of learning is also possible in the paleo-cortex of the human brain, as well as other non-motor, non-sensory regions. These structures are largely dissociated from the dorsal striatal system.[14] However, despite this dissociation, we demonstrated that participants learned to modulate activity in a specific and volitional manner, much like other neocortical regions explored in non-human primates and rodents. Motor skill learning, and neuroprosthetic skill learning, proceeds in a prototypical manner, where the early phase of learning is characterized by a rapid acquisition of task parameters, following by a slower refinement process.[10–12,36] The experimental sessions in this study were not long enough to investigate the later stages of learning, but we robustly demonstrate the early stage of learning, characterized by rapid changes in the firing characteristics of the direct neurons. While limbic structures do not directly project to the dorsal striatum, they do project heavily to the ventral striatum. In fact, the ventral striatum is thought to serve as the interface between the limbic and motor systems.[24,63] So, is it possible that the type of neuroprosthetic skill learning that we demonstrate here is facilitated by the ventral striatum instead?

While we cannot answer this question by directly recording activity from the ventral striatum in humans, we sought out neural signatures of this interaction. Neely and colleagues[10] previously demonstrated that neuroprosthetic skill learning in the primary visual cortex increasingly recruited the striatum. With knockout experiments, they further demonstrated that the dorsomedial striatum (which receives direct projections from V1) was essential for the acquisition of such neuroprosthetic skill. In their work, a prominent neural signature of this striatal involvement was increased SFC between the spiking activity of the direct neuron and LFPs in the 10–25 Hz band. In our work, we observed a similar increase in SFC in the high alpha/low beta (10–20 Hz) bands as learning progressed. This increase in SFC was independent of power or firing rate changes, was specific to the direct neurons and occurred only in learning sessions. Since oscillations in this frequency band are rarely observed in the medial temporal lobe (where delta, theta and gamma oscillations dominate[42,64–66]) this observed SFC is probably the result of an external influence on the local hippocampal oscillations. In rodents, Lansink and colleagues[45] demonstrated that beta oscillations in the hippocampus can be driven by reward-predictive cues, and enhanced by learning. They also demonstrated that hippocampal spiking activity can be phase-locked to the underlying beta oscillations by reward-predictive cues. Furthermore, they demonstrated that learning can also increase SFC between neurons in the ventral striatum and beta oscillations in the hippocampus. Thus, beta oscillations in the hippocampus and related structures may be driven by a reward prediction mechanism, potentially driven by the ventral tegmental afferents to CA1,[67] or indirectly from the striatum via the ventral pallidal-mediodorsal thalamic route.[68] Furthermore, the observed SFC is unlikely related to attentional modulation, since attention has actually been shown to decrease the alpha/beta band SFC in the visual cortex.[69] Taken together, the learning-specific SFC observed in this study probably implicates the ventral striatum in

this type of neuroprosthetic skill learning in mnemonic structures. Although this argument is largely speculative, anatomical connectivity of limbic structures to the ventral striatum provides additional supporting evidence.

One of the canonical characteristics of the cortico-basal ganglia loops is the presence of parallel inhibitory and disinhibitory pathways,[60,70] which allow the basal ganglia to play a role in selection of context relevant motor plans or even cognitive strategies.[61] The medium spiny neurons (MSNs), which are ubiquitous within the basal ganglia, are furnished with dopamine receptors in close proximity to the corticostriatal terminals.[71] Dopaminergic innervation of these MSNs by the midbrain dopaminergic system facilitates plastic synaptic changes that shapes the striatal and resulting basal ganglia outputs, playing a role in facilitating reward-based learning. Interestingly, the ventral striatum is also known to form cortico-basal ganglia loops,[70] with a variety of limbic structures and the anterior cingulate cortex as its primary input and output.[14,72–74] Since the MSNs that compose much of the striatum are difficult to excite,[60] convergent input from the limbic structures and the anterior cingulate cortex could drive ventral striatal MSNs, activating a series of parallel inhibitory and disinhibitory circuits that can be actively tuned via the midbrain dopaminergic system to facilitate reward-based learning of precise limbic activity patterns. Future work in animal models will certainly focus on interrogating this limbic-basal ganglia circuitry to establish the significance of the ventral striatum in facilitating this type of limbic neuroprosthetic skill learning.

It is important to contrast the work presented here from the recent developments in the field of BMIs. In traditional BMIs, the machine learns to decode an intent from a particular set of neural activity, and then generates a control signal to perform a particular behaviour. Due to significant technical advancements over the past two decades, the complexity of neural activity that we can record and decode information from continues to grow, leading to remarkable clinical possibilities such as the ability to control complex robotic arms,[75] reanimate limbs[75,76] and even decode handwriting with a high degree of accuracy.[77] Compared to this, modulation of a single neuron to control one-dimensional movement of a block on a screen may seem paltry. However, there is a major conceptual difference between the approach taken in this paper and the traditional field of BMIs. Whereas traditional BMIs are concerned with maximizing decoding performance, our intention is to use BMIs to interrogate basic neuroscience questions. In our approach, which was pioneered by Fetz and colleagues, we emphasize the brain, rather than the machine, as the primary learning entity. The algorithm that maps the selected neural activity to the defined behaviour is extremely simple and static, but the onus is on the conscious brain to learn to modulate itself to achieve a particular behaviour. BMIs are unique in the sense that they enforce a genuinely causal, rather than merely correlative link between a particular neural activity and the corresponding behaviour. This allows us to study old concepts of learning and plasticity with new methodology. This is even more true in our case, where the activity to be modulated resides in mnemonic structures. In our experiments, the brain as an observer learns to modulate activity in these structures in a likely striatal/basal ganglia dependent manner. However, these structures themselves are involved in a variety of mnemonic and learning-related processes. Hence, our work opens up a brand-new avenue of questioning: how do these traditionally distinct learning systems (i.e. the basal ganglia and the medial temporal lobe) interact? Can they be forced to cooperate or compete? What effect does this interaction have on downstream behaviour? These are the types of question that we can begin to address with the type of methodology presented here.

The data presented here indicate that single neuron activity in limbic and other memory-related structures can be precisely regulated in a rapid, highly specific and volitional manner in humans. Furthermore, this type of neuroprosthetic skill learning in limbic structures is probably facilitated by the limbic-basal ganglia circuity involving the ventral striatum. Such, high-fidelity self-regulation of neural activity may provide an avenue for the development of novel neuroprosthetics for the treatment of neurological conditions that commonly present with pathological activity in limbic structures, such as medically refractory epilepsy. Furthermore, since limbic structures, and particularly those of the medial temporal lobe, are critical to mnemonic processes, obtaining volitional control over highly specific activity in these structures may provide a mechanism of probing the function and plasticity of these brain structures without exogenous stimulation.

Acknowledgements

We would like to acknowledge Victoria Barkley for her help with organizing and assisting with patient data collection, Andrea Gómez Palacio for her help with patient data collection, Ryan Ramos for his contributions towards developing the neurofeedback task and David Groppe for his advice on choosing appropriate statistical analyses.

Funding

This work was funded by the Natural Sciences and Engineering Research Council (RGPIN-2015–05936 to T.A.V. and RGPIN-2016–06358 to M.R.P), National Institutes of Health and Brain Canada (grant number 1U01NS103792-01 to T.A.V.), Canadian Fund for Innovation and Ontario Research Fund (grant number 35923), Vanier Canada Graduate Scholarships (to K.P. and C.N.K), Toronto General and Western Foundation, Dean Connor and Maris Uffelmann Donation and the Walter & Maria Schroeder Institute.

Competing interests

M.R.P. is a Director and Co-Founder of the company MyndTec Inc., which has no involvement in this study. T.A.V. is an investor in the company Neurescence, which has no involvement in this study. T.A.V. provides consulting to the company Panaxium, which has no involvement in this study.

Supplementary material

Supplementary material is available at *Brain* online.

References

1. Lozano AM, Fosdick L, Chakravarty MM, et al. A phase II study of fornix deep brain stimulation in mild Alzheimer's disease. *J Alzheimer's Dis*. 2016;54(2):777–787.
2. Lee DJ, Lozano AM. Current status of deep brain stimulation for Alzheimer's disease: From chance observation to clinical trials. *Cold Spring Harb Symp Quant Biol*. 2018;83:201–205.
3. Kuhn J, Hardenacke K, Lenartz D, et al. Deep brain stimulation of the nucleus basalis of Meynert in Alzheimer's dementia. *Mol Psychiatry*. 2015;20(3):353–360.
4. Hescham S, Lim LW, Jahanshahi A, et al. Deep brain stimulation of the forniceal area enhances memory functions in experimental dementia: The role of stimulation parameters. *Brain Stimul*. 2013;6(1):72–77.

5. Ezzyat Y, Wanda PA, Levy DF, et al. Closed-loop stimulation of temporal cortex rescues functional networks and improves memory. *Nat Commun*. 2018;9(1):365.
6. Fetz EE, Baker MA. Operantly conditioned patterns on precentral unit activity and correlated responses in adjacent cells and contralateral muscles. *J Neurophysiol*. 1973;36(2):179–204.
7. Fetz EE. Operant conditioning of cortical unit activity. *Science*. 1969;163(3870):955–958.
8. Prsa M, Galiñanes GL, Huber D. Rapid integration of artificial sensory feedback during operant conditioning of motor cortex neurons. *Neuron*. 2017;93(4):929–939.e6.
9. Clancy KB, Koralek AC, Costa RM, Feldman DE, Carmena JM. Volitional modulation of optically recorded calcium signals during neuroprosthetic learning. *Nat Neurosci*. 2014;17(6):807–809.
10. Neely RM, Koralek AC, Athalye VR, Costa RM, Carmena JM. Volitional modulation of primary visual cortex activity requires the basal ganglia. *Neuron*. 2018;97(6):1356–1368.e4.
11. Koralek AC, Costa RM, Carmena JM. Temporally precise cell-specific coherence develops in corticostriatal networks during learning. *Neuron*. 2013;79(5):865–872.
12. Koralek AC, Jin X, Long JD, et al. Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature*. 2012;483(7389):331–335.
13. Kemp JM, Powell TPS. The cortico-striate projection in the monkey. *Brain*. 1970;93(3):525–546.
14. Voorn P, Vanderschuren LJMJ, Groenewegen HJ, Robbins TW, Pennartz CM. Putting a spin on the dorsal-ventral divide of the striatum. *Trends Neurosci*. 2004;27(8):468–474.
15. Lubar JF, Bahler WW. Behavioral management of epileptic seizures following EEG biofeedback training of the sensorimotor rhythm. *Biofeedback Self Regul*. 1976;1(1):77–104.
16. Lubar JF, Shabsin HS, Natelson SE, et al. EEG operant conditioning in intractable epileptics. *Arch Neurol*. 1981;38(11):700–704.
17. Ros T, Baars B, Lanius RA, Vuilleumier P. Tuning pathological brain oscillations with neurofeedback: A systems neuroscience framework. *Front Hum Neurosci*. 2014;8:1008.
18. Sitaram R, Ros T, Stoeckel L, et al. Closed-loop brain training: The science of neurofeedback. *Nat Rev Neurosci*. 2017;18(2):86–100.
19. Walker JE. Using QEEG-guided neurofeedback for epilepsy versus standardized protocols: Enhanced effectiveness? *Appl Psychophysiol Biofeedback*. 2010;35(1):29–30.
20. Corlier J, Rimsky-Robert D, Valderrama M, et al. Self-induced intracerebral gamma oscillations in the human cortex. *Brain*. 2016;139(Pt 12):3084–3091.
21. Corlier J, Valderrama M, Navarrete M, et al. Voluntary control of intracortical oscillations for reconfiguration of network activity. *Sci Rep*. 2016;6:36255.
22. Cerf M, Thiruvengadam N, Mormann F, et al. On-line, voluntary control of human temporal lobe neurons. *Nature*. 2010;467(7319):1104–1108.
23. Aflalo T, Kellis S, Klaes C, et al. Decoding motor imagery from the posterior parietal cortex of a tetraplegic human. *Science*. 2015;348(6237):906–910.
24. Berke JD, Okatan M, Skurski J, Eichenbaum HB. Oscillatory entrainment of striatal neurons in freely moving rats. *Neuron*. 2004;43(6):883–896.
25. Groppe DM, Bickel S, Dykstra AR, et al. iELVis: An open source MATLAB toolbox for localizing and visualizing human intracranial electrode data. *J Neurosci Methods*. 2017;281:40–48.
26. Rutishauser U, Schuman EM, Mamelak AN. Online detection and sorting of extracellularly recorded action potentials in human medial temporal lobe recordings, in vivo. *J Neurosci Methods*. 2006;154(1-2):204–224.
27. Rutishauser U, Schuman EM, Mamelak AN. Online detection and sorting of extracellularly recorded action potentials in human medial temporal lobe recordings, in vivo. *J Neurosci Methods*. 2006;154(1-2):204–224.
28. Faraut MCM, Carlson AA, Sullivan S, et al. Data descriptor: Dataset of human medial temporal lobe single neuron activity during declarative memory encoding and recognition. *Sci Data*. 2018;5:180010.
29. Kober SE, Witte M, Ninaus M, Neuper C, Wood G. Learning to modulate one's own brain activity: The effect of spontaneous mental strategies. *Front Hum Neurosci*. 2013;7:695.
30. Garcia-Garcia MG, Marquez-Chin C, Popovic MR. Operant conditioning of motor cortex neurons reveals neuron-subtype-specific responses in a brain-machine interface task. *Sci Rep*. 2020;10(1):19992.
31. Gourévitch B, Eggermont JJ. A nonparametric approach for detection of bursts in spike trains. *J Neurosci Methods*. 2007;160(2):349–358.
32. Gregoriou GG, Gotts SJ, Zhou H, Desimone R. High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science*. 2009;324(5931):1207–1210.
33. Hammer EM, Halder S, Blankertz B, et al. Psychological predictors of SMR-BCI performance. *Biol Psychol*. 2012;89(1):80–86.
34. Athalye VR, Santos FJ, Carmena JM, Costa RM. Evidence for a neural law of effect. *Science*. 2018;359(6379):1024–1029.
35. Fries P. Rhythms for cognition: Communication through coherence. *Neuron*. 2015;88(1):220–235.
36. Athalye VR, Carmena JM, Costa RM. Neural reinforcement: Re-entering and refining neural dynamics leading to desirable outcomes. *Curr Opin Neurobiol*. 2020;60:145–154.
37. Rutishauser U, Ross IB, Mamelak AN, Schuman EM. Human memory strength is predicted by theta-frequency phase-locking of single neurons. *Nature*. 2010;464(7290):903–907.
38. Rizzuto DS, Madsen JR, Bromfield EB, et al. Reset of human neocortical oscillations during a working memory task. *Proc Natl Acad Sci U S A*. 2003;100(13):7931–7936.
39. Lega BC, Jacobs J, Kahana M. Human hippocampal theta oscillations and the formation of episodic memories. *Hippocampus*. 2012;22(4):748–761.
40. Guderian S, Schott BH, Richardson-Klavehn A, Düzel E. Medial temporal theta state before an event predicts episodic encoding success in humans. *Proc Natl Acad Sci U S A*. 2009;106(13):5365–5370.
41. Cohen MX. *Analyzing neural time series data*. MIT Press; 2014.
42. Katz CN, Patel K, Talakoub O, Groppe D, Hoffman K, Valiante TA. Differential generation of saccade, fixation, and image-onset event-related potentials in the human mesial temporal lobe. *Cereb Cortex*. 2020;30(10):5502–5516.
43. Groppe DM, Bickel S, Keller CJ, et al. Dominant frequencies of resting human brain activity as measured by the electrocorticogram. *Neuroimage*. 2013;79:223–233.
44. Lalla L, Rueda Orozco PE, Jurado-Parras MT, Brovelli A, Robbe D. Local or not local: Investigating the nature of striatal theta oscillations in behaving rats. *eNeuro*. 2017;4(5):ENEURO.0128-17.2017.
45. Lansink CS, Meijer GT, Lankelma JV, Vinck MA, Jackson JC, Pennartz CMA. Reward expectancy strengthens CA1 theta and beta band synchronization and hippocampal-ventral striatal coupling. *J Neurosci*. 2016;36(41):10598–10610.
46. Arduin P-J, Frégnac Y, Shulz DE, Ego-Stengel V. Ego-Stengel V. Bidirectional control of a one-dimensional robotic actuator by operant conditioning of a single unit in rat motor cortex. *Front Neurosci*. 2014;8(8):206.
47. Eaton RW, Libey T, Fetz EE. Operant conditioning of neural activity in freely behaving monkeys with intracranial reinforcement. *J Neurophysiol*. 2017;117(3):1112–1125.

48. Kobayashi S, Schultz W, Sakagami M. Operant conditioning of primate prefrontal neurons. *J Neurophysiol*. 2010;103(4): 1843–1855.
49. Moritz CT, Fetz EE. Volitional control of single cortical neurons in a brain-machine interface. *J Neural Eng*. 2011;8(2):025017.
50. Clancy KB, Koralek AC, Costa RM, Feldman DE, Carmena JM. Volitional modulation of optically recorded calcium signals during neuroprosthetic learning. *Nat Neurosci*. 2014;17(6): 807–809.
51. Arduin P-J, Fregnac Y, Shulz DE, Ego-Stengel V. 'Master' neurons induced by operant conditioning in rat motor cortex during a brain-machine interface task. *J Neurosci*. 2013;33(19):8308–8320.
52. Houweling AR, Brecht M. Behavioural report of single neuron stimulation in somatosensory cortex. *Nature*. 2008;451(7174): 65–68.
53. Sadtler PT, Quick KM, Golub MD, et al. Neural constraints on learning. *Nature*. 2014;512(7515):423–426.
54. D'Avella A. Modularity for motor control and motor learning. In: Laczko J, Latash ML, eds. *Progress in motor control*. Springer International Publishing; 2017:3–19.
55. Wulf G, Lewthwaite R. Optimizing performance through intrinsic motivation and attention for learning: The OPTIMAL theory of motor learning. *Psychon Bull Rev*. 2016;23(5):1382–1414.
56. Curcio G, Ferrara M, De Gennaro L. Sleep loss, learning capacity and academic performance. *Sleep Med Rev*. 2006;10(5):323–337.
57. Greeley B, Seidler RD. Mood induction effects on motor sequence learning and stop signal reaction time. *Exp Brain Res*. 2017;235(1):41–56.
58. Van Wouwe NC, Band GPH, Ridderinkhof KR. Positive affect modulates flexibility and evaluative control. *J Cogn Neurosci*. 2011;23(3):524–539.
59. Best MD, Takahashi K, Suminski AJ, Ethier C, Miller LE, Hatsopoulos NG. Comparing offline decoding performance in physiologically defined neuronal classes. *J Neural Eng*. 2016; 13(2):026004.
60. Groenewegen HJ. The basal ganglia and motor control. *Neural Plast*. 2003;10(1-2):107–120.
61. Graybiel AM. The basal ganglia and cognitive pattern generators. *Schizophr Bull*. 1997;23(3):459–469.
62. McGeorge AJ, Faull RLM. The organization of the projection from the cerebral cortex to the striatum in the rat. *Neuroscience*. 1989;29(3):503–537.
63. Mogenson GJ, Jones DL, Yim CY. From motivation to action: Functional interface between the limbic system and the motor system. *Prog Neurobiol*. 1980;14(2-3):69–97.
64. Jutras MJ, Fries P, Buffalo EA. Oscillatory activity in the monkey hippocampus during visual exploration and memory formation. *Proc Natl Acad Sci U S A*. 2013;110(32):13144–13149.
65. Jutras MJ, Fries P, Buffalo EA. Gamma-band synchronization in the macaque hippocampus and memory formation. *J Neurosci*. 2090;29(40):12521–12531.
66. Montefusco-Siegmund R, Leonard TK, Hoffman KL. Hippocampal gamma-band synchrony and pupillary responses index memory during visual search. *Hippocampus*. 2017;27(4): 425–434.
67. Groenewegen HJ, Berendse HW, Haber SN. Organization of the output of th
e ventral striatopallidal system in the rat: Ventral pallidal efferents. *Neuroscience*. 1993;57(1):113–142.
68. Groenewegen HJ, Wright CI, Uylings HBM. The anatomical relationships of the prefrontal cortex with limbic structures and the basal ganglia. *J Psychopharmacol*. 1997;11(2):99–106.
69. Buffalo EA, Fries P, Landman R, Buschman TJ, Desimone R. Laminar differences in gamma and alpha coherence in the ventral stream. *Proc Natl Acad Sci U S A*. 2011;108(27):11262–11267.
70. Alexander GE, Delong MR, Strick PL. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annu Rev Neurosci*. 1986;9(1):357–381.
71. Smith AD, Bolam JP. The neural network of the basal ganglia is revealed by the study of synaptic connections of identified neurones. *TINS*. 1990;13(7):259–265.
72. Graybiel AM. Building action repertoires: Memory and learning functions of the basal ganglia. *Curr Opin Neurobiol*. 1995;5(6): 733–741.
73. Kunishio K, Haber SN. Primate cingulostriatal projection: Limbic striatal versus sensorimotor striatal input. *J Comp Neurol*. 1994;350(3):337–356.
74. Thierry A-M, Gioanni Y, Dégénétais E, Glowinski, J. Hippocampo-prefrontal cortex pathway: Anatomical and electrophysiological characteristics. *Hippocampus*. 2000;10(4): 411–419.
75. Ajiboye AB, Willett FR, Young DR, et al. Restoration of reaching and grasping movements through brain-controlled muscle stimulation in a person with tetraplegia: A proof-of-concept demonstration. *Lancet*. 2017;389(10081):1821–1830.
76. Moritz CT, Perlmutter SI, Fetz EE. Direct control of paralysed muscles by cortical neurons. *Nature*. 2008;456(7222):639–642.
77. Willett FR, Avansino DT, Hochberg LR, Henderson JM, Shenoy KV. High-performance brain-to-text communication via handwriting. *Nature*. 2021;593(7858):249–254.