



Learning in a closed-loop brain-machine interface with distributed optogenetic cortical feedback

Dorian Goueytes, Henri Lassagne, Daniel Shulz, Valérie Ego-Stengel, Luc Estebanez

► To cite this version:

Dorian Goueytes, Henri Lassagne, Daniel Shulz, Valérie Ego-Stengel, Luc Estebanez. Learning in a closed-loop brain-machine interface with distributed optogenetic cortical feedback. *Journal of Neural Engineering*, 2022, 19 (6), pp.066045. 10.1088/1741-2552/acab87 . hal-03915216v2

HAL Id: hal-03915216

<https://hal.science/hal-03915216v2>

Submitted on 29 Dec 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Learning in a closed-loop brain-machine interface with distributed optogenetic cortical feedback

Dorian Goueytes¹, Henri Lassagne¹, Daniel E. Shulz¹, Valérie Ego-Stengel^{1†} and Luc Estebanez^{1†*}

¹Université Paris-Saclay, CNRS, Institut de Neurosciences Paris-Saclay, 91400 Saclay, France.

*Corresponding author: Email: luc.estebanez@cnrs.fr

† These authors contributed equally.

Abstract

Objective Distributed microstimulations at the cortical surface can efficiently deliver feedback to a subject during the manipulation of a prosthesis through a brain-machine interface. Such feedback can convey vast amounts of information to the prosthesis user and may be key to obtain an accurate control and embodiment of the prosthesis. However, so far little is known of the physiological constraints on the decoding of such patterns. Here, we aimed to test a rotary optogenetic feedback that was designed to encode efficiently the 360° movements of the robotic actuators used in prosthetics. We sought to assess its use by mice that controlled a prosthesis joint through a closed-loop brain-machine interface.

Approach We tested the ability of mice to optimize the trajectory of a virtual prosthesis joint in order to solve a rewarded reaching task. They could control the speed of the joint by modulating the activity of individual neurons in the primary motor cortex. During the task, the patterned optogenetic stimulation projected on the primary somatosensory cortex continuously delivered information to the mouse about the position of the joint.

Main results We showed that mice are able to exploit the continuous, rotating cortical feedback in the active behaving context of the task. Mice achieved better control than in the absence of feedback by detecting reward opportunities more often, and also by moving the joint faster towards the reward angular zone, and by maintaining it longer in the reward zone. Mice controlling acceleration rather than speed of the joint failed to improve motor control.

Significance These findings suggest that in the context of a closed-loop brain-machine interface, distributed cortical feedback with optimized shapes and topology can be exploited to control movement. Our study has direct applications on the closed-loop control of rotary joints that are frequently encountered in robotic prostheses.

1. Introduction

Invasive motor brain-machine interfaces (BMIs) focus on establishing a stable translation from brain neuronal activity into motor commands (Carmena et al., 2003; Chapin et al., 1999)(Collinger et al., 2013; Wodlinger et al., 2015). Most implementations rely on visual feedback to guide the prosthesis during the task. In particular, they lack somatosensory feedback like touch and proprioception. This feedback is however critical for movement accuracy, as shown by studies in humans in which local peripheral anesthesia blocking afferent tactile sensation reduced dexterity and impaired fine motor control of the hand (Johansson and Westling, 1984; Monzée et al., 2003)(Flesher et al., 2021).

Implementing somatosensory-like feedback from a prosthesis back to the subject requires first to fit the prosthesis with touch and proprioceptive-like sensors (Roberts et al., 2021), and then to relay this information to the central nervous system. In invasive closed-loop brain-machine interfaces, feedback stimulation generally targets the primary somatosensory cortex, where neuronal activation is integrated as touch inputs in the awake behaving rodent (O'Connor et al., 2013; Prsa et al., 2017; Sachidhanandam et al., 2013) and primate (O'Doherty et al., 2011) as well as in humans (Flesher et al., 2021).

Strategies to provide behaviorally-relevant input using such cortical stimulation often rely on the intensity or frequency modulation of a stimulation targeting one spatially limited region of interest, which limits the amount of information that can be delivered (O'Doherty et al., 2019, 2011; Prsa et al., 2017). However, recent technical progress has made distributed neuronal activations possible, by using sophisticated multichannel electrical microstimulations (Dadarlat et al., 2015; Fernández et al., 2021; Flesher et al., 2021; Weiss et al., 2019) or by harnessing spatiotemporally patterned optogenetic stimulation of the cortex (Abbasi et al., 2018; Ceballo et al., 2019; Goueytes et al., 2019; Lassagne et al., 2022). Such distributed neuronal activation at the surface of the cortex can convey multiple information streams in parallel (Hartmann et al., 2016), such as those arising from the multiple touch-like sensors (Roberts et al., 2021) that are available in modern bidirectional prostheses (D'Anna et al., 2019; Johannes et al., 2011). Further, feedback spread across a large cortical area can be more robust and provides an opportunity to mimic physiological, distributed peripheral inputs by generating spatiotemporal patterns of activation that embrace the well-known topography of primary sensory areas, including the primary somatosensory cortex (Abbasi et al., 2018; Flesher et al., 2021, 2016), the primary visual cortex (Chen et al., 2020; Dobelle et al., 1976; Dobelle, 2000) and the auditory cortex (Ceballo et al., 2019).

This emerging capability to integrate distributed cortical feedback in a brain-machine interface (Pandarinath and Bensmaia, 2022) raises multiple questions. One is that most robotic prostheses are fitted with rotary actuators that each drive one degree of freedom. The instantaneous angular position of these actuators is a critical information that should be channeled back to the subject. But such circular information cannot be conveyed unambiguously by modulating the activity in a single spot of the cortex. In earlier work (Lassagne et al., 2022), we explored the use of a spatially distributed, continuous spatio-temporal pattern of photoactivation to convey this information in the form of a rotating bar projected on the surface of the somatosensory cortex of awake behaving mice. However, this previous study was a purely passive sensory task, while during physiological behavior, sensory integration cannot be dissociated from active motor control (Poulet and Petersen, 2008).

Here, based on a recently developed closed-loop brain-machine interface (Abbasi et al., 2018; Goueytes et al., 2019) we asked if this feedback could be efficiently exploited by the mice beyond sensory processing, by helping to control the angular position of a simulated prosthesis.

We found that the mice were able not only to detect the location of the Reward zone, but also to alter the dynamics of the rotary joint. Notably, they learned to increase the speed of the movement towards the Reward zone while preserving the amount of time spent in it. This was not observed in trials where the feedback photostimulation was not activated.

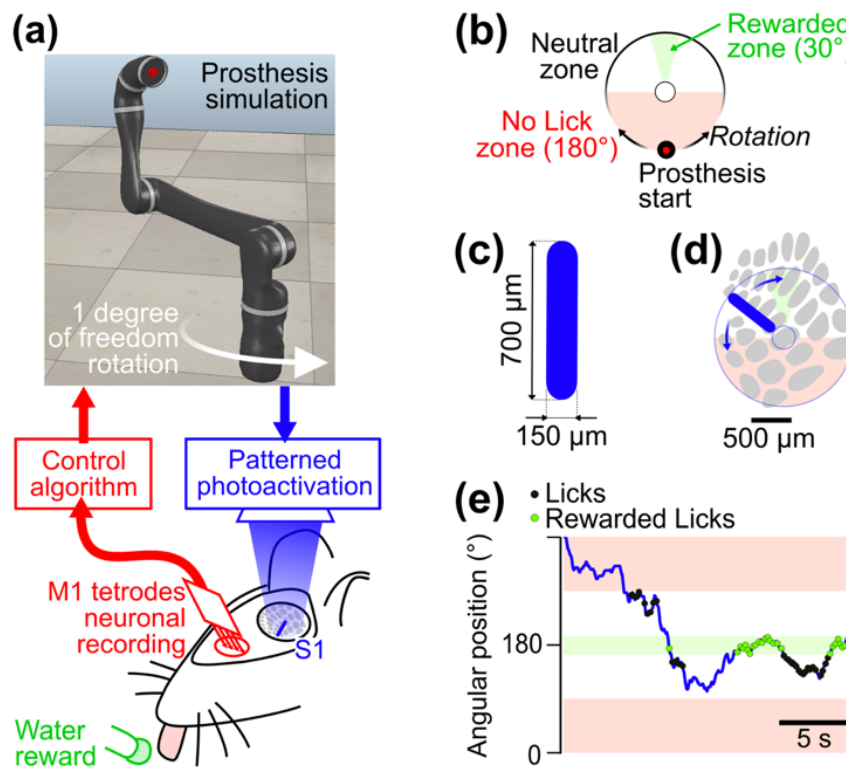


Figure 1 Closed-loop control of a virtual prosthesis with patterned optogenetic somatosensory feedback (a) Schematic representation of the closed-loop setup, including a snapshot of the prosthesis simulation in V-REP. The most proximal joint of the prosthesis is controlled by neuronal activity recorded in M1. The angular position of the joint is fed back to the mouse through spatio-temporally patterned photoactivation of S1.

(b) The motor space explored by the prosthesis during the task was divided in distinct functional zones. Licks occurring when the prosthesis was in the Reward zone triggered water rewards. Licks occurring when the prosthesis was in the No lick zone aborted the trial. The Start position of the prosthesis was always located at 0°.

(c) Shape of the photostimulation pattern projected on the cortex.

(d) Location of the targeted cortical surface with respect to the barrels of the whisker primary somatosensory cortex. The photostimulation bar rotates in synchrony with the proximal joint in prosthesis space in (b). The Start position of the photostimulation corresponded to the most posterior position of the stimuli in brain coordinates.

(e) Example trajectory of the rotating prosthesis, and associated licks (black dots) and rewarded licks (green dots), during the 20 s of one uninterrupted trial.

2. Methods

We developed a neuroprosthetic bidirectional brain-machine interface in mice by combining electrophysiological recordings in M1 with patterned optogenetic stimulations in S1, using mice expressing channelrhodopsin in excitatory cortical neurons (see Methods, (Goueytes et al., 2019)). Mice had to control a single rotary joint of an off-the-shelf prosthesis (Jaco 2, Kinova Robotics) that was simulated using the V-REP software (Rohmer et al., 2013). The mice were trained to perform an exploratory sensorimotor task in this 360° angular space without boundaries, using as the sole feedback a dynamic patterned photostimulation projected on the surface of the barrel cortex (figure 1(a)).

2.1 Surgical preparation

We report data from 13 *Emx1-Cre;Ai27* mice (both male and female) expressing channelrhodopsin in excitatory neurons across the cortex (Madisen et al., 2012). All animals were implanted via two successive surgeries under Isoflurane anesthesia (1-4% Isoflurane mixed in 100% air), combined with local Lidocaine analgesia (< 7mg/kg) and anti-inflammatory drug Meloxicam (1-8 mg/kg). During the first surgery, the cranial skin was resected, the skull was exposed, and after careful removal of residual soft tissue and complete air drying, a 5 mm diameter craniotomy was drilled over the “barrel” whisker area of primary somatosensory cortex (S1, - 1.5 mm P, -3.3 mm L relative to Bregma). At the same time, we labelled the position of the whisker area of the primary motor cortex (M1, +1.5mm P, -0.6mm L relative to Bregma).

A 5 mm diameter glass window was then positioned in direct contact with the dura mater, and sealed with cyanoacrylate glue (Loctite “Super Glue Power Gel”) to the skull, thereby resulting in a chronic, stable optical access to the barrel cortex. The remaining exposed skull was coated with a layer of liquid cyanoacrylate glue. A head-fixation bar was also glued on the skull, contralateral to the optical window side. Low reflection, black dental cement (Ortho Jet, Lang dental, USA) was applied on the skull to protect it and anchor the bar and the window. After a week of recovery, we mapped the barrel cortex through the optical window using intrinsic imaging (Lassagne et al., 2022). In particular, we identified the location of the barrel corresponding to the C2 whisker. In a second surgery, using a similar preparation as the first one, we went back to the previously identified location of whisker M1; we drilled and opened the skull, removed the dura mater locally, and implanted in a chronic fashion (Okun et al., 2016) an extracellular recording electrode (32 channel silicon probe, A1x32Poly35mm25s 177A32, Neuronexus, USA). We descended the tip 800 µm below the surface of the cortex, thereby targeting neurons from Layer 5 of M1.

2.2 Neuronal electrophysiological recordings

After electrode implantation, we monitored electrophysiological signals daily to control the stability of recording. After about one week, we started to acquire neuronal recordings. We manually isolated large amplitude units based on waveform shape and cluster separation (Blackrock microsystems, USA). All neuronal signals were sampled at 30 kHz and the records were stored. A detailed description of the quality of the signals and of their stability across sessions is available in (Abbasi et al., 2019), where the very same methods were used. Given the duration of the experiments described here, we estimated that a single neuron picked

arbitrarily at the start of the training may not be recorded stably across the whole training period. This was the motivation behind the choice of training a set of 6-8 master neurons during the experiment, so that the recording of a subset of these neurons would always be stable while training continued.

2.3 Prosthetic simulation

A commercial Jaco2 (Kinova robotics) prosthesis was simulated using the robotic simulation software V-REP (Rohmer et al., 2013). This simulation was based on a CAD model provided by the manufacturer, and cross-validated during the loan of a physical Jaco2 unit.

To connect the virtual robotic arm to our brain-machine interface (figure 1(a)), we developed a custom V-REP plugin. Thanks to this software bridge, the speed and direction commands for the proximal joint of the virtual prosthesis was updated approximately every 12 ms, based on the neuronal activity readout. In return, the current angular position of the joint was fed back to the brain-machine interface and was used to update the angular position of the optogenetic S1 stimulation (Goueytes et al., 2019).

We measured an end-to-end latency of the combined BMI software and robotic arm system of 36 ms (standard deviation 4 ms).

2.4 Control algorithm

At the beginning of the first training session, among the neurons that were manually spike sorted, two arbitrary groups of 3 to 4 Master neurons were selected using an automatic algorithm that minimized the difference in population firing rate between the two groups.

For 7 mice, speed control of the joint was implemented by linearly translating the population firing rate of each group into a speed command. To compute the speed command in one group, the activity of all neurons in that group was summed together and sampled every 10 ms, followed by convolution with a 100 ms box kernel. The speed command resulting from the first (resp. second) group was assigned to the clockwise (resp. counterclockwise) direction. The difference between the two speeds was directly relayed as a speed command to the proximal joint of the V-REP model of the prosthesis, leading it to rotate in its 360° circular space (figure 1(b)), without any limit to the extent of the circular rotations. Note that in speed control mode, the Kinova arm acceleration was set to a maximum of approximately 150°/s².

To calibrate the linear relationship between the smoothed group firing rate and the corresponding speed command, we computed the average activity of the two groups in a 5 minutes “baseline” waiting period at the beginning of each training session, and the firing rate of each group was divided by this baseline value. This normalization ensured that the velocity distribution was centered around 0°/s for each session, despite day-to-day variability in individual firing rates. Finally, we scaled the resulting speed by a constant selected during the first session, defined so that the average absolute baseline joint speed would be ~30°/s. The resulting distributions of angular speeds in the first and last sessions are shown in figure 5(c).

In the last part of our study, we switched from speed to acceleration control on a separate group of 6 mice. We applied the exact same computation of neuronal activity, but the final conversion was into an acceleration command rather than a speed command. We adjusted the linear factor in the conversion so that the observed distributions of position and speed would be as close as possible to those observed in the speed control condition (figure 5(b,c)).

2.5 Optogenetic stimulation patterns

We designed optogenetic feedback patterns that took the form of a 700 μm long and 150 μm wide bar centered on the C2 barrel (figure 1(c,d)), based on the results of the intrinsic imaging session. The photostimulation was generated with a Digital Light Processing module (DLP, Vialux V-7001, Germany) containing a 1024 x 768 pixels Texas Instruments micro-mirror chip coupled with a high-power 462 nm blue LED. During training, the rotation of the bar over the somatosensory cortex followed at short latency the position of the arm. The C2 barrel was excluded from the stimulation in order to avoid its overstimulation. We have previously provided a full description of the illumination optics and validation experiments (Abbasi et al., 2018).

2.6 Task and behavioral training

Once the mice were implanted with the electrodes and the BMI appeared functional, they were water regulated in order to enable operant conditioning based on water reward (figure 1(a)). Their weight was monitored and maintained at 80% of their baseline (measured prior to the start of the water deprivation) by supplementing water if needed. Simultaneous to the start of the water regulation, the mice were habituated to head-fixation and trained to obtain rewards by licking a reward port placed in front of them (1-2 sessions). After habituation, training was performed on a daily basis without interruption for the whole duration of the training.

Each training session corresponded to approximately 10 min of online spike sorting, 5 min of baseline recording, 40 min of behavioral training and 5 min of post-training recording, for a total of one hour per day. The sessions were divided into trials lasting at most 20 s. At the beginning of each trial, the rotary joint position was initiated at 0°. The mice could freely displace the joint across a circular space in both directions by modulating the activity of the neurons in the two Master groups (figure 1(b)). The mice were rewarded only if they licked while the joint was located in a 30° zone centered on the 180° position.

In order to discourage strategies based on continuous licking independently from the joint position, we aborted any trial where licks occurred while the arm was in a 180° zone around the starting position (the “No lick zone”). This was followed by a 2 s timeout. The intertrial interval lasted for a minimum of 5 s and a maximum of 7 s in case of a timeout following a lick in the No lick zone. To maintain interest in the task over the full session, we averted long series of rewards by interrupting trials after more than 4 s continuously in the Reward zone. The optogenetic feedback was turned on at the onset of each trial, and switched off at its offset. An example trial is shown in figure 1(e).

In order to specifically test the contribution of the feedback to the mouse behavior, in 4 animals, at all phases of training, we included 20% of trials interleaved randomly where no optogenetic feedback was provided.

Training was interrupted when the number of neurons recorded in M1 became too low for the task requirement (less than 6), or when optical access to S1 through the glass window was degraded (opaque, blurry window).

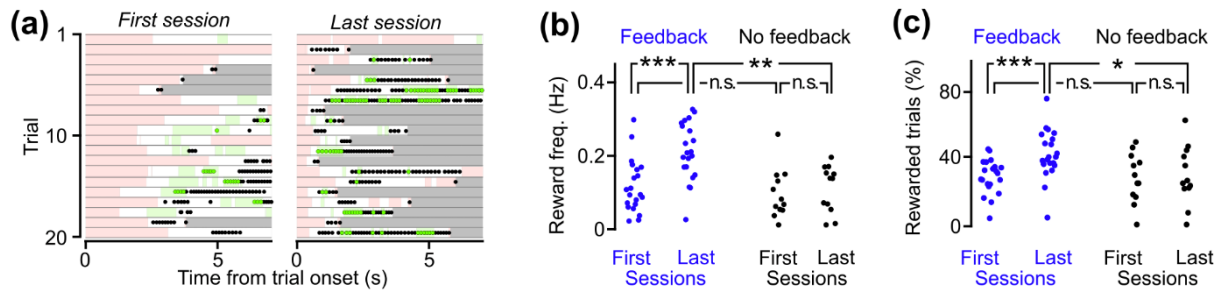


Figure 2 Feedback enables increased task performance after training.

(a) Temporal sequence of spatial zones explored by the prosthesis during the first 20 trials of the first session (left) and of the last session (right) for which the joint reached at least once the Neutral zone. Each line represents the time course of a trial. Pink: No lick zone, in which licks interrupt the trial. Green: Reward zone. White: Neutral zone. Gray: the trial has already stopped. Black dots: licks. Green dots: rewarded licks. Note that licks take place in bursts, defined as lick sequences during which inter-lick interval is at most 300 ms.

(b) Reward frequency averaged across each session, plotted for the first three and the last three training sessions of each mouse. Blue: Optogenetic feedback (7 mice, 21 sessions). Black: Control, no feedback (subset of 4 mice, 12 sessions).

(c) Same as b, for the proportion of rewarded trials, calculated as the percentage of trials where at least one lick was rewarded. Trials where the joint never went past the “No lick” zone were excluded.

All statistical comparisons are obtained from Mann-Whitney tests. *: $p < 0.05$ **: $p < 0.01$. ***: $p < 0.001$. Detailed statistics for each test, including number of samples and exact p value, are provided in the main text.

3. Results

We trained 7 mice to obtain rewards by manipulating the speed of a virtual rotating joint through the control of the neuronal activity of motor cortex neurons. Mice received feedback about the angular position of the joint by optogenetic stimulation in the somatosensory cortex (see Methods, figure 1). As a control, in 4 of the 7 mice no optogenetic feedback was available during 20% of trials. Mice were trained daily for at least 17 consecutive sessions (average 17.8 sessions). Training was interrupted when recording quality degraded or technical difficulties arose with the recordings. The last training session was systematically excluded from further analysis.

3.1 Mice learn to improve their performance in the feedback condition

Across training, the mice learned to increase the number of rewards they collected in the feedback condition. This can be observed on the example of figure 2(a) showing the motor exploration and licking behavior during the first 20 trials of the first and last session for one mouse. At the population level, the reward frequency was significantly larger in the last three training sessions compared to the first three sessions, when feedback was available (figure 2(b) left, Mann-Whitney $p = 0.00024$, $N = 21$ sessions from 7 mice, null hypothesis: “there is no increase in reward frequency with learning”), but not in the no-feedback condition (figure 2(b) right, $p = 0.33$, $N = 12$ sessions from 4 mice). At the end of training, the mice performance in trials with feedback was also significantly higher than in trials when no feedback was provided (figure 2(b), $p = 0.0016$; feedback: $N = 21$; no-feedback: $N = 12$, null hypothesis: “there is no difference in the reward frequency between the feedback and no-feedback condition”), in contrast to the start of the training ($p = 0.38$; feedback: $N = 21$; no-feedback: $N = 12$).

The role of optogenetic feedback in enabling learning was also clear when computing the proportion of trials that were rewarded. Trials where optogenetic feedback was available resulted in more rewarded trials in the last three training sessions compared to the first three sessions (figure 2(c) left, Mann-Whitney $p = 0.00082$, $N = 21$ sessions), while this was not the case in the no-feedback condition (figure 2(c) right, $p = 0.93$, $N = 12$ sessions). Finally, after training, trials with optogenetic feedback were significantly more rewarded than no-feedback trials (figure 2(c), $p = 0.041$; feedback: $N = 21$; no feedback: $N = 12$). This was not the case at the start of the training ($p = 0.96$; feedback: $N = 21$; no feedback: $N = 12$). Note that in this analysis of percentages of rewarded trials, only trials reaching the No lick border were included.

3.2 Mice learn to detect and anticipate entry in the Reward zone

To better characterize the mice ability to track the joint position and lick appropriately, we first measured the proportion of rewarded trials among those where the joint entered the Reward zone. We found that in the optogenetic feedback condition, the proportion of detected entries in the Reward zone increased significantly between the first and last three training sessions (figure 3(a) left, Mann-Whitney $p = 0.0030$, $N = 21$ sessions), but not in the no-feedback condition (figure 3(a) right, $p = 0.52$, $N = 12$ sessions). As a result, at the end of training, the proportion of rewardable trials was significantly larger with optogenetic feedback than without (figure 3(a), $p = 0.0011$; feedback: $N = 21$; no feedback: $N = 12$).

We then looked at the delay from the entry of the joint in the Reward zone to the first rewarded lick. We found that with optogenetic feedback, the mice learned to lick faster following entry in the Reward zone (figure 3(b) left, Mann-Whitney $p = 6.3 \times 10^{-7}$, $N = 21$ sessions). This was not the case without optogenetic feedback (figure 3(b) right, $p = 0.32$, $N = 12$ sessions). Note that after training, the difference between the delay to first lick measured in the feedback versus

no-feedback conditions was not quite significant (figure 3(b), $p = 0.056$; feedback: $N = 21$; no feedback: $N = 12$).

Finally, we found that the mice learned to collect rewards at a significantly higher frequency once inside the Reward zone (figure 3(c) left, Mann-Whitney $p = 3.3 \times 10^{-6}$, $N = 21$ sessions). This increase was not visible in the absence of feedback (figure 3(c) right, $p = 0.35$, $N = 12$ sessions), and after training, the frequency of licking in the reward zone was significantly higher with feedback compared to the no-feedback condition (figure 3(c), $p = 1.1 \times 10^{-4}$; feedback: $N = 21$; no feedback: $N = 12$). Overall, these results suggest that the mice did manage to exploit the optogenetic feedback to track the joint position and obtain rewards.

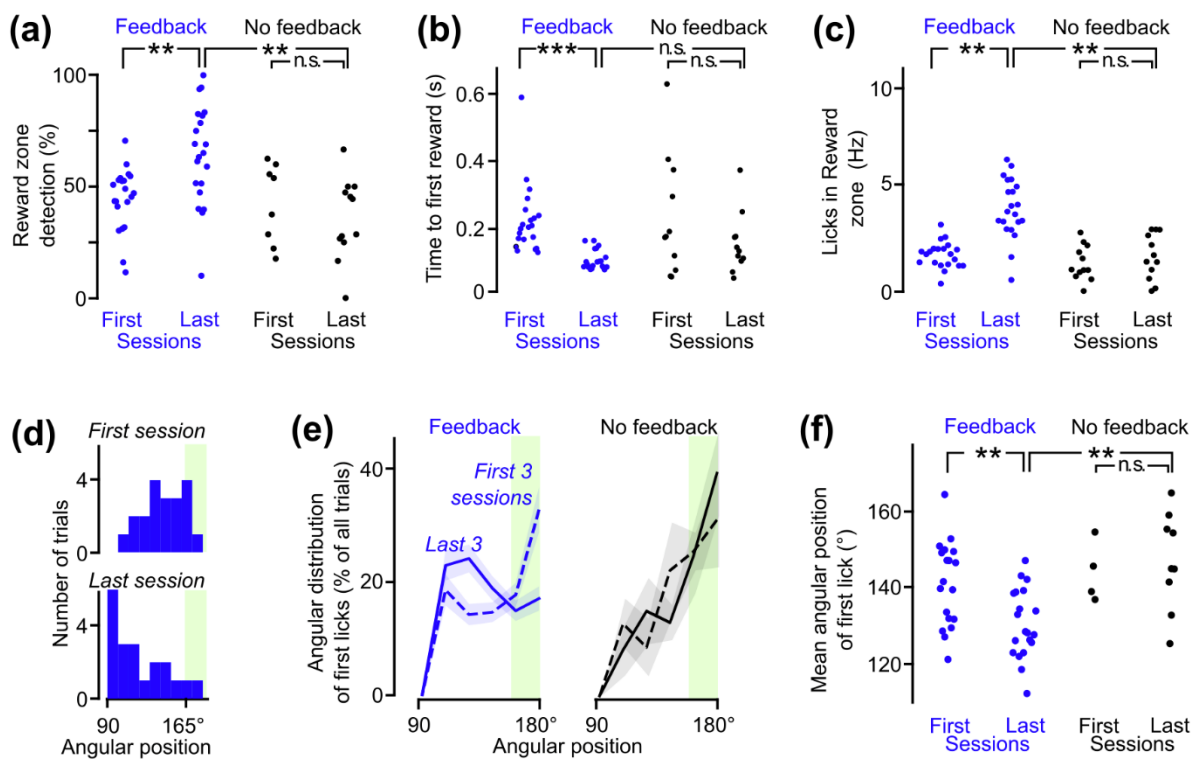


Figure 3 Feedback enables accurate spatial anticipation and detection of the Reward zone.

(a) Percentage of rewarded trials among those where the joint reached the Reward zone. Blue: Optogenetic feedback (7 mice, 21 sessions). Black: Control, no feedback (subset of 4 mice, 12 sessions). (b) Mean delay from entry in the Reward zone to the first rewarded lick.

(c) Mean licking frequency when the joint was inside the Reward zone.

(d) Distribution of the angular position of the first lick of each trial in one mouse, during the first and the last training session, with the optogenetic feedback activated. In these graphs, angular positions larger than 180° were mirrored to represent all data within a 0 - 180° rather than 0 - 360° range, given the left-right symmetry of our circular space. Color code as in Figure 1. Trials that were interrupted by a first lick in the No lick zone are excluded.

(e) Population distribution of the angular position of the first lick (see case study in d). Light background: standard error of the mean (SEM) computed over all sessions of each group.

(f) Average angular position of the first lick for each session, computed from the same dataset as in (e). Detailed statistics for each test, including number of samples and exact p values, are provided in the main text.

To better understand how the mice took advantage of the spatial organization of the optogenetic feedback, we focused on the licking patterns that occurred around the entry in the Reward zone, and analyzed the corresponding spatial position of the joint. In this analysis, we focused on trials that led to rewards, because we aimed to characterize the behavioral sequences that are successful. Note that as a consequence, none of the trials that started (and ended) with a lick in the ‘No Lick zone’ were included in the analysis. We noticed that in the feedback condition, the onset of licking bursts took place for increasingly early angles after the transition between the No lick and the Neutral zones (example first and last sessions for one mouse in figure 3(d)). This shift was visible across the population on the histogram of the mean angular position of the first lick after trial onset (figure 3(e)). Consistently, the average angular position of the first lick significantly diminished between the first three and last three training sessions in the feedback condition (figure 3(f) left, Mann-Whitney $p = 0.0021$, $N = 21$ sessions), but not in the no-feedback condition (figure 3(f) right, $p = 0.71$, $N = 12$ sessions), and at the end of training, the average first lick angular position was significantly lower in the feedback condition than in the no-feedback condition (figure 3(f), $p = 0.004$, feedback: $N = 21$; no feedback: $N = 12$). This anticipation of licking towards the border of the No lick zone in the feedback condition did not result in an increase in premature licks in the No lick zone, that would have canceled the trial. Instead, there was a significant reduction in aborted trials during learning with feedback (on average from 76% down to 64% of all initiated trials, Mann-Whitney $p = 0.42$, $N = 21$ sessions).

We hypothesized that this anticipation strategy may be driven by the mice tendency to perform long bouts of repeated licks (“bursts”, where the interval between licks was of less than 300 ms) rather than individual licks. Indeed, generating a long lick burst before entering the Reward zone may be an efficient way to ensure that multiple licks land inside the Reward zone. Consistent with this hypothesis, we found that the mice generated increasingly long bursts of licks when they exited for the first time the No lick zone in a trial (from 6.8 to 16.3 licks per burst, first three vs. last three sessions, Mann-Whitney $p = 6.0 \times 10^{-6}$, $N = 21$ sessions). These first bursts in the trials were increasingly long enough to include licks that occurred while the robotic joint was located in the Reward zone, and therefore the proportion of trials where the first burst included rewarded licks increased on average from 6 to 15% during training (Mann-Whitney $p = 0.0013$, $N = 21$ sessions).

Overall, the mice learned that they could start licking as soon as the prosthesis joint entered the Neutral zone, and they took full advantage of this to initiate long bursts of licking that enabled them to collect multiple rewards.

3.3 Mice learn to accelerate movements of the prosthesis to the Reward zone, and stabilize it using feedback.

So far, our analysis showed that during this task, the continuous, rotative bar-like optogenetic feedback provided sufficient information for the mice to increase their reward rate with training, by generating timely licking bursts. Beyond the contribution of feedback processing to performance, we next asked if the mice were able to actively modulate the rotation of the joint towards the Reward zone. To quantify motor control, we first focused on the beginning of the trial, before the first entry in the Reward zone. We found that the mice learned to spend less time in the Neutral zone as they moved the joint towards the Reward zone (case studies in figure 4(a)). This reduction was significant when feedback was available (figure 4(b) left, Mann-Whitney $p = 0.013$, $N = 21$ sessions) but non-significant without feedback (figure 4(b) right, $p = 0.65$, $N = 12$ sessions). Note that after training, we found no significant difference between the feedback and no-feedback conditions (figure 4(b), $p = 0.25$, feedback: $N = 21$; No feedback: $N = 12$).

This faster movement towards the Reward zone stemmed from an increased average angular speed of the joint before entry in the Reward zone. We observed this increased speed in the feedback condition (computed over 50 ms bins, figure 4(c) left, Mann-Whitney $p = 0.0015$, $N = 21$ sessions) but not in the no-feedback condition (figure 4(c) right, $p = 0.088$, $N = 12$ sessions). Note that after training, the difference in speed measured in the feedback versus no-feedback conditions was not significant (figure 4(c), $p = 0.055$; feedback: $N = 21$; no feedback: $N = 12$).

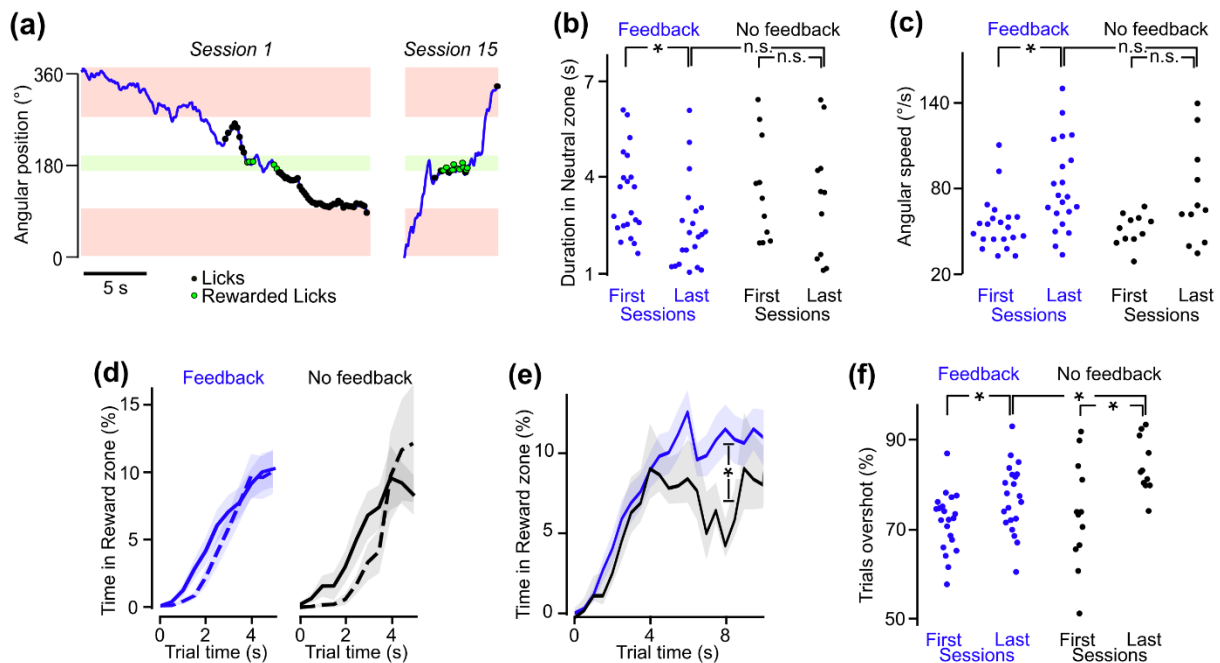


Figure 4 Mice move the rotary joint faster to the Reward zone with training, and can stabilize better in the Reward zone with feedback.

(a) Two example trials from the same mouse, in the first versus the last training session.

(b) Duration spent in the Neutral zone before entering for the first time in the Reward zone, in the first three versus the last three training sessions. Each point is a different mouse/session, same dataset and conventions as in Figures 2 and 3.

(c) Average angular speed of the rotary joint between trial onset and the first entry in the Reward zone.

(d) Percentage of time spent in the Reward zone over the first five seconds of the trials (bins: 0.5 s). Left: Feedback sessions. Right: No-feedback sessions. Dashed lines: trials from the first 3 sessions. Continuous lines: trials from the last 3 sessions. Light background: SEM across sessions.

(e) Percentage of time spent in the Reward zone in the feedback versus the no-feedback condition after training, over the first 10 s of the trials. Blue: feedback. Black: no feedback. Light background: SEM across sessions.

(f) Percentage of trials where, after entering the Neutral zone and then the Reward zone, the robot joint overshoot and reached the No lick zone on the other side of the circular space.

Detailed statistics for each test, including number of samples and exact p value, are provided in the main text.

These changes in joint movements after learning resulted in longer time spent in the Reward zone in the first seconds of the trials, regardless of the availability of feedback (figure 4(d)). However, the presence of feedback did impact the rotary joint motor control later in the trials. When we looked at the time course up to 10 seconds in the trial, we noticed that the percentage of time in the Reward zone increased and then, after 5 s, stabilized above 10% in the feedback condition, whereas it dropped back below 5% in the no-feedback condition (figure 4(e)). Note however that this difference was only significant in one 500 ms interval, 8 s after onset (Mann-Whitney $p = 0.0099$, feedback: $N = 21$ sessions; no feedback: $N = 12$ sessions). Given that we tested for significance the 5 intervals where the SEM of the two conditions did not overlap, the Bonferroni correction led to a significance threshold of 0.01 that was just met by the p -value of the test.

The difference in average time in the Reward zone led us to hypothesize that in the feedback trials, the mice were able to exploit the photostimulation to stabilize the joint in or close to the Reward zone, and thus increase the percentage of time spent there. By contrast, in the absence of feedback, the mice could not adjust the trajectory of the joint as well, possibly not decelerating early enough because of the lack of feedback. To test this hypothesis, we measured the proportion of trials where, after entering the Neutral zone from one side of the circular space, the rotary joint would go on, overshoot the Reward and Neutral zones, and exit the Neutral zone on the other side. We found that both with and without feedback, the proportion of overshoot trials increased significantly after training (figure 4(f); feedback: Mann-Whitney $p = 0.032$, $N = 21$ sessions; no-feedback: $p = 0.012$, $N = 12$ sessions). This was probably due to the faster joint movements, which led to many trajectories leaving the Reward zone fast. However, we found that this proportion was significantly smaller when the feedback photostimulation was available (figure 4(f), $p = 0.010$; feedback: $N = 21$; no feedback: $N = 12$), so that the mouse managed more often to stabilize the joint in the Reward zone or in the Neutral zone.

Overall, these observations suggest that mice learned to control the rotary joint movement. Better control was achieved with optogenetic feedback of the joint angular position, by allowing trajectory adjustments during the trial.

3.4 Motor control of the rotary joint depends on the controlled variable.

The fact that the mice learned to move the rotary joint faster, and moreover that they could adjust dynamically their motor control of the joint while the trials with feedback were ongoing, indicated that a motor control algorithm based on an angular speed command was successful. This prompted us to ask whether other motor control algorithms could be used. We hypothesized that controlling the angular acceleration of a rotary joint may be more straightforward, because it is directly related to muscle torque and thus possibly to neural activity. To test this, we trained 6 additional mice in the same closed-loop setting, but with an acceleration neural controller (example trajectory in figure 5(a)). Like in the speed condition, control of the rotary joint was achieved by the modulation of the firing rate of two antagonist groups of 3 to 4 neurons each (see Methods).




Figure 5 Mice fail to control the prosthesis movements using an acceleration-based joint controller.

(a) Example trajectory of the prosthesis in the angular space, in a mouse trained with a controller that translated M1 firing rates into acceleration, instead of the speed of the rotary joint.

(b) Distributions of the angular position of the rotary prosthesis joint in mice trained with the acceleration-based joint controller (purple, 6 mice) versus the speed-based controller (blue, 7 mice). Dashed lines: first 3 sessions. Continuous lines: last 3 sessions. The session start position is always 0°. Distributions are averaged across sessions/mice. Light background: standard error of the mean.

(c) Distributions of the angular speed of the rotary joint for the same sessions as in (b).

(d) Reward frequency averaged across each session, plotted for the first three and the last three training sessions of each mouse for the acceleration-based controller.

(e) Percentage of trials where the joint entered the Reward zone for which the mouse obtained a reward, computed across each session/mouse, for the first three versus last three training sessions. The three tested conditions are shown: Acceleration control, always with optogenetic feedback; Speed control, with feedback, and Speed control, without feedback. The speed data is the same as Figure 3(a).

(f) Average angular position of the first lick of rewarded bursts for the same sessions as in (e). The speed data is the same as Figure 3(f).

(g) Average instantaneous angular speed of the rotary joint between trial onset and the first entry in the Reward zone for the same sessions as in (e). The speed data is the same as Figure 4(c).

(h) Percentage of time spent in the Reward zone over the first ten seconds of the acceleration trials (bins: 0.5 s). Dashed lines: trials from the first 3 sessions. Continuous lines: trials from the last 3 sessions.

Detailed statistics, including number of samples and exact p values, are provided in the main text.

The average distributions of the position and speed of the rotary joint in the two control algorithm conditions were in the same ranges (figure 5(b,c)). Importantly, during the first training sessions, reward opportunities were as frequent in the two conditions. Mice spent 4.3% of the total time in the Reward zone in the speed control, versus 4.0% of the time in the acceleration control condition (not significantly different, Mann-Whitney $p = 1.0$; speed control: $N = 21$ sessions from 7 mice; acceleration control: $N = 18$ sessions from 6 mice). In addition, on average the joint moved with comparable speed through the circular space ($47^\circ/\text{s}$ vs $34^\circ/\text{s}$, not significantly different, Mann-Whitney $p = 0.44$). These results suggest that the photostimulation patterns conveying information about the joint angular position activated the cortex in similar ways.

The 6 mice that were trained with the acceleration controller received significantly more rewards in the last sessions of training compared to the first sessions, from a reward frequency of 0.066 rewards per second up to 0.17 (figure 5(d), Mann-Whitney $p = 0.0054$, $N = 18$ sessions). The mice also learned to lick for reward more often when the prosthesis trajectory reached the Reward zone (first vs. last sessions, figure 5(e), Mann-Whitney $p = 0.0012$, $N = 18$ sessions from 6 mice). In addition, like in the speed-controller condition, the mice anticipated the entry of the prosthesis into the Reward zone, and started licking before it happened (figure 5(f), Mann-Whitney $p = 0.007$, $N=18$). These results suggest that these mice learned to exploit the feedback information to lick at appropriate times. However, we failed to find any sign of active motor control or of motor learning. In particular, none of the changes in joint dynamics that we observed in the speed control condition were present in the acceleration control condition. For example, the average speed of the controller (figure 5(g), Mann-Whitney $p = 0.42$, $N=18$), the delay to enter the Reward zone, and the time spent in the Reward zone (figure 5(h)) were not modified through learning with the acceleration controller. We conclude from these data that the mice were unable to control the prosthesis movements through the acceleration-based controller, while they could still benefit from the feedback about its angular position.

4. Discussion

Using a closed-loop brain-machine interface controlling a prosthesis simulation in the mouse model, we have shown that a mesoscopic scale feedback with a circular topography and rotating dynamics can be exploited to optimize the motor control of a robotic joint. Specifically, the mice learned to improve both the rotation of the joint and the detection of the target location in order to obtain more rewards.

4.1 Integration of the optogenetic stimulation in open versus closed-loop experiments

Our experiments show that the guidance of licking based on the decoding of a mesoscopic cortical feedback can be as successful during an active, motor brain-machine interface task, as during a purely passive detection task (Lassagne et al., 2022). We therefore conclude that the sensory feedback part of a closed-loop brain-machine interface can be efficiently designed and tested independently from the motor part of the loop. This is consistent with the recent successful development of an invasive closed-loop neuroprosthesis (Flesher et al., 2021), that was based on the independent validation of the microstimulations in the primary somatosensory cortex of humans (Flesher et al., 2016) and primates (Kim et al., 2015).

4.2 Anticipated licking

We found that throughout training, the mice learned to initiate their licking bursts increasingly early, and ultimately as soon as the photostimulation entered the Neutral zone, where licks were

not rewarded but did not interrupt the trials (figure 3). It was also mainly in the Neutral zone that the mice increased the speed of the rotary joint as it headed towards the Reward zone. Together, these findings suggest that the mice motivation to obtain rewards was high enough to initiate licking as soon as possible without canceling the trial. At the same time, maintaining uninterrupted licking bursts was costly enough to lead the mice to increase the speed of the joint as it headed towards the Reward zone.

4.3 Limited motor learning

Overall, the two changes in motor strategy during learning with the speed-based motor control were an increased speed of the rotary joint at the beginning of the trials, followed by later adjustments to preserve the time spent in the Reward zone by limiting overshooting (figure 4(f)). We hypothesize that the limited scale of these changes was due to the design of the task, since in the naive state before training, the mice already spent approximately 6% of the trial time in the Reward zone (see Results section 3.4, and figure 4(d,e) for a time course comparison). In future experiments, it would be interesting to reduce the size of the Reward zone, in order to encourage more active stabilization of the joint in the Reward zone.

4.4 Comparison of speed versus acceleration-based controllers

The experiments based on controlling acceleration rather than speed highlighted the key impact of controller design on closed-loop performance and motor control (figure 5(g,h)). We are not aware of previous works describing BMIs based on acceleration. We hypothesize that the relative inefficiency of the acceleration-based controller may be due to the challenge of connecting the acceleration values to the position feedback, given that there are no less than two steps of temporal integration between acceleration and position, versus only one in the case of the speed controller. Nonetheless, these experiments confirmed that mice can take advantage of the sensory feedback regardless of their ability to control the actuator (figure 5(e,f)).

4.5 Related works

To our knowledge, this is the first demonstration of a motor BMI that includes a distributed, dynamical cortical feedback that takes into account the mechanical constraints of a joint. Here, we aimed at optimizing the feedback based on a robotic implementation of a joint. We first tested the perception of cortical stimulation patterns in a passive detection task (Lassagne et al., 2022). We specifically designed a 360° rotating pattern to encode the angle of the joint, similar to a proprioceptive signal, while still abiding to the continuity of the representations that can be found in the barrel cortex that we targeted. Our experimental results suggest that a similar approach may address the challenge of proprioceptive-like feedback in human neuroprosthetics (Delhayé et al., 2018).

In our work, we chose to implement distributed optogenetics in order to write feedback information directly in the cortex. This approach has been used before in open-loop experiments in order to explore the functional properties of primary sensory cortices (Ceballo et al., 2018; Lassagne et al., 2022). In a BMI task, it was pioneered by (Prsa et al., 2017) in the rodent. Translating it to the non-human primates and to humans, poses a specific set of challenges (Shen et al., 2020). Therefore, so far, distributed stimulations of primary sensory cortex in primates and humans have been achieved through lower-resolution, invasive arrays of microelectrodes (Dadarlat et al., 2015; Flesher et al., 2021, 2016; Kim et al., 2015).

Focusing more specifically on the control and behavioral paradigm, we can relate our study to a line of work initiated by Carmena and collaborators (Clancy et al., 2014; Koralek et al., 2012).

These experiments are based on continuous feedback, in that case in auditory space, and on the assumption that neuronal plasticity enables BMI learning. In addition, similar to our work, continuous control of a cursor was achieved by rodents through the modulation of the difference between the firing rate of two ensembles of M1 neurons.

However, one key difference with our work is that the neuronal activity, motor control and feedback spaces were linearly related. The firing rate of neurons controlled directly the position of the cursor from a minimum to a maximum position, and feedback was provided in the form of a continuous sound with a pitch that varied also linearly with the cursor position.

In contrast, in our work the cursor space was circular, and the cursor could rotate without boundaries, as it was controlled by modulating its speed or its acceleration.

It appears that motor learning was stronger and faster in those studies (Clancy et al., 2014; Koralek et al., 2012) than in our work. As we hypothesized above, this may be mainly due to the fact that in these studies, both the feedback and control were tied to the same variable (position). In contrast, in our study the feedback was tied to position in a more complex, circular space, while control was related to its first or second derivative (speed/acceleration). This distance between the control and feedback variables may have made learning and control more difficult in our setting.

5. Conclusion

Our findings suggest that using patterned cortical feedback with a spatio-temporal structure relevant to physical constraints is an efficient strategy to provide critical information about the position of a prosthesis, and support its active control. Here we have tested this strategy to encode the angular position of a single rotary joint. It could be extended to explore the simultaneous encoding of the angular position of multiple joints, as well as the addition of tactile-like feedback of strategic points on the prosthesis surface. Future experiments will be needed to probe the ability of mice to integrate such multiple-dimensional feedbacks to actively control a full prosthesis.

Data availability statement

The data that support the findings of this study are available upon request from the authors.

Acknowledgment

We thank Aurélie Daret for help with experiments and lab managing. We thank Dr. Maria Makarov at CentraleSupélec for support regarding the robot modeling and control. We thank the robotic company Sysaxes (France) for the loan of a Kinova Jaco2 unit. We thank the following funding sources: ANR Neurowhisk ANR-14-CE24-0019; ANR Mesobrain ANR-20-CE37-0013; ANR Motorsense ANR-21-CE37-0012; FRM DEQ20170336761; iCODE and IDEX Paris-Saclay ANR-11-IDEX-0003-02; FRC AAP2018; CNRS 80 Prime 2020, and Fondation Dassault Systèmes.

Conflict of interest

The authors declare no competing interests.

Ethics approval and consent to participate

All animal experiments were performed according to European and French law as well as CNRS guidelines and were approved by the French ministry for research (ethical committee 59, authorization 858-2015060516116339v5).

Data and code availability statement

The data and code that support the findings of this study are available from the corresponding authors, upon reasonable request.

Authorship contribution statement

DES, VE and LE initiated and supervised the project. DG, DES, VE and LE designed the experiments. DG performed the experiments. LE, DG and HL analyzed the data. LE wrote the paper with inputs from all authors.

References

- Abbasi, A., Estebanez, L., Goueytes, D., Lassagne, H., Shulz, D.E., Ego-Stengel, V., 2019. Cortical closed-loop brain-machine interface requires biomimetic sensory feedback (preprint). *BioRxiv*. <https://doi.org/10.1101/2019.12.12.873794>
- Abbasi, A., Goueytes, D., Shulz, D.E., Ego-Stengel, V., Estebanez, L., 2018. A fast intracortical brain-machine interface with patterned optogenetic feedback. *J. Neural Eng.* 15, 046011. <https://doi.org/10.1088/1741-2552/aabb80>
- Carmena, J.M., Lebedev, M.A., Crist, R.E., O'Doherty, J.E., Santucci, D.M., Dimitrov, D.F., Patil, P.G., Henriquez, C.S., Nicolelis, M.A.L., 2003. Learning to Control a Brain-Machine Interface for Reaching and Grasping by Primates. *PLoS Biol.* 1, e42. <https://doi.org/10.1371/journal.pbio.0000042>
- Ceballo, S., Bourg, J., Kempf, A., Piwkowska, Z., Daret, A., Deneux, T., Rumpel, S., Bathellier, B., 2018. Cortical recruitment determines learning dynamics and strategy. *bioRxiv*. <https://doi.org/10.1101/274936>
- Ceballo, S., Piwkowska, Z., Bourg, J., Daret, A., Bathellier, B., 2019. Targeted Cortical Manipulation of Auditory Perception. *Neuron* S0896627319308463. <https://doi.org/10.1016/j.neuron.2019.09.043>
- Chapin, J.K., Moxon, K.A., Markowitz, R.S., Nicolelis, M.A.L., 1999. Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nat. Neurosci.* 2, 664–670. <https://doi.org/10.1038/10223>
- Chen, X., Wang, F., Fernandez, E., Roelfsema, P.R., 2020. Shape perception via a high-channel-count neuroprosthesis in monkey visual cortex. *Science* 370, 1191–1196. <https://doi.org/10.1126/science.abd7435>
- Clancy, K.B., Koralek, A.C., Costa, R.M., Feldman, D.E., Carmena, J.M., 2014. Volitional modulation of optically recorded calcium signals during neuroprosthetic learning. *Nat. Neurosci.* 17, 807–809. <https://doi.org/10.1038/nn.3712>
- Collinger, J.L., Wodlinger, B., Downey, J.E., Wang, W., Tyler-Kabara, E.C., Weber, D.J., McMorland, A.J., Velliste, M., Boninger, M.L., Schwartz, A.B., 2013. High-performance neuroprosthetic control by an individual with tetraplegia. *The Lancet* 381, 557–564. [https://doi.org/10.1016/S0140-6736\(12\)61816-9](https://doi.org/10.1016/S0140-6736(12)61816-9)
- Dadarlat, M.C., O'Doherty, J.E., Sabes, P.N., 2015. A learning-based approach to artificial sensory feedback leads to optimal integration. *Nat. Neurosci.* 18, 138–144. <https://doi.org/10.1038/nn.3883>
- D'Anna, E., Valle, G., Mazzoni, A., Strauss, I., Iberite, F., Patton, J., Petrini, F.M., Raspopovic, S., Granata, G., Di Iorio, R., 2019. A closed-loop hand prosthesis with simultaneous intraneural tactile and position feedback. *Sci. Robot.* 4, eaau8892.
- Delhay, B.P., Long, K.H., Bensmaia, S.J., 2018. Neural Basis of Touch and Proprioception in Primate Cortex, in: Terjung, R. (Ed.), *Comprehensive Physiology*. Wiley, pp. 1575–1602. <https://doi.org/10.1002/cphy.c170033>
- Dobelle, W.H., Mladejovsky, M.G., Evans, J.R., Roberts, T.S., Girvin, J.P., 1976. 'Braille' reading by a blind volunteer by visual cortex stimulation. *Nature* 259, 111–112. <https://doi.org/10.1038/259111a0>
- Dobelle, Wm.H., 2000. Artificial Vision for the Blind by Connecting a Television Camera to the Visual Cortex: *ASAIJ* 46, 3–9. <https://doi.org/10.1097/00002480-200001000-00002>
- Fernández, E., Alfaro, A., Soto-Sánchez, C., Gonzalez-Lopez, P., Lozano, A.M., Peña, S., Grima, M.D., Rodil, A., Gómez, B., Chen, X., Roelfsema, P.R., Rolston, J.D., Davis, T.S., Normann, R.A., 2021. Visual percepts evoked with an intracortical 96-channel microelectrode array inserted in human occipital cortex. *J. Clin. Invest.* 131, e151331. <https://doi.org/10.1172/JCI151331>
- Flesher, S.N., Collinger, J.L., Foldes, S.T., Weiss, J.M., Downey, J.E., Tyler-Kabara, E.C., Bensmaia, S.J., Schwartz, A.B., Boninger, M.L., Gaunt, R.A., 2016. Intracortical microstimulation of human somatosensory cortex. *Sci. Transl. Med.* 8, 361ra141–361ra141. <https://doi.org/10.1126/scitranslmed.aaf8083>

- Flesher, S.N., Downey, J.E., Weiss, J.M., Hughes, C.L., Herrera, A.J., Tyler-Kabara, E.C., Boninger, M.L., Collinger, J.L., Gaunt, R.A., 2021. A brain-computer interface that evokes tactile sensations improves robotic arm control. *Science* 372, 831–836. <https://doi.org/10.1126/science.abd0380>
- Goueytes, D., Abbasi, A., Lassagne, H., Shulz, D.E., Estebanez, L., Ego-Stengel, V., 2019. Control of a robotic prosthesis simulation by a closed-loop intracortical brain-machine interface, in: 2019 9th International IEEE/EMBS Conference on Neural Engineering (NER). Presented at the 2019 9th International IEEE/EMBS Conference on Neural Engineering (NER), IEEE, San Francisco, CA, USA, pp. 183–186. <https://doi.org/10.1109/NER.2019.8716885>
- Hartmann, K., Thomson, E.E., Zea, I., Yun, R., Mullen, P., Canarick, J., Huh, A., Nicolelis, M.A.L., 2016. Embedding a Panoramic Representation of Infrared Light in the Adult Rat Somatosensory Cortex through a Sensory Neuroprosthesis. *J. Neurosci.* 36, 2406–2424. <https://doi.org/10.1523/JNEUROSCI.3285-15.2016>
- Johannes, M.S., Bigelow, J.D., Burck, J.M., Harshbarger, S.D., Kozlowski, M.V., Doren, T.V., 2011. An Overview of the Developmental Process for the Modular Prosthetic Limb. *JOHNS HOPKINS APL Tech. Dig.* 30, 10.
- Johansson, R.S., Westling, G., 1984. Roles of glabrous skin receptors and sensorimotor memory in automatic control of precision grip when lifting rougher or more slippery objects. *Exp. Brain Res.* 56. <https://doi.org/10.1007/BF00237997>
- Kim, S., Callier, T., Tabot, G.A., Gaunt, R.A., Tenore, F.V., Bensmaia, S.J., 2015. Behavioral assessment of sensitivity to intracortical microstimulation of primate somatosensory cortex. *Proc. Natl. Acad. Sci.* 112, 15202–15207. <https://doi.org/10.1073/pnas.1509265112>
- Koralek, A.C., Jin, X., Long II, J.D., Costa, R.M., Carmena, J.M., 2012. Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature* 483, 331–335. <https://doi.org/10.1038/nature10845>
- Lassagne, H., Goueytes, D., Shulz, D.E., Estebanez, L., Ego-Stengel, V., 2022. Continuity within the somatosensory cortical map facilitates learning. *Cell Rep.* 39, 110617. <https://doi.org/10.1016/j.celrep.2022.110617>
- Madisen, L., Mao, T., Koch, H., Zhuo, J., Berenyi, A., Fujisawa, S., Hsu, Y.-W.A., Garcia, A.J., Gu, X., Zanella, S., Kidney, J., Gu, H., Mao, Y., Hooks, B.M., Boyden, E.S., Buzsáki, G., Ramirez, J.M., Jones, A.R., Svoboda, K., Han, X., Turner, E.E., Zeng, H., 2012. A toolbox of Cre-dependent optogenetic transgenic mice for light-induced activation and silencing. *Nat. Neurosci.* 15, 793–802. <https://doi.org/10.1038/nn.3078>
- Monzée, J., Lamarre, Y., Smith, A.M., 2003. The Effects of Digital Anesthesia on Force Control Using a Precision Grip. *J. Neurophysiol.* 89, 672–683. <https://doi.org/10.1152/jn.00434.2001>
- O'Connor, D.H., Hires, S.A., Guo, Z.V., Li, N., Yu, J., Sun, Q.-Q., Huber, D., Svoboda, K., 2013. Neural coding during active somatosensation revealed using illusory touch. *Nat. Neurosci.* 16, 958–965. <https://doi.org/10.1038/nn.3419>
- O'Doherty, J.E., Lebedev, M.A., Ifft, P.J., Zhuang, K.Z., Shokur, S., Bleuler, H., Nicolelis, M.A.L., 2011. Active tactile exploration using a brain-machine-brain interface. *Nature* 479, 228–231. <https://doi.org/10.1038/nature10489>
- O'Doherty, J.E., Shokur, S., Medina, L.E., Lebedev, M.A., Nicolelis, M.A.L., 2019. Creating a neuroprosthesis for active tactile exploration of textures. *Proc. Natl. Acad. Sci.* 116, 21821–21827. <https://doi.org/10.1073/pnas.1908008116>
- Okun, M., Lak, A., Carandini, M., Harris, K.D., 2016. Long Term Recordings with Immobile Silicon Probes in the Mouse Cortex. *PLOS ONE* 11, e0151180. <https://doi.org/10.1371/journal.pone.0151180>
- Pandarathna, C., Bensmaia, S.J., 2022. The science and engineering behind sensitized brain-controlled bionic hands. *Physiol. Rev.* 102, 551–604. <https://doi.org/10.1152/physrev.00034.2020>
- Poulet, J.F.A., Petersen, C.C.H., 2008. Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* 454, 881–885. <https://doi.org/10.1038/nature07150>
- Prsa, M., Galiñanes, G.L., Huber, D., 2017. Rapid Integration of Artificial Sensory Feedback during Operant Conditioning of Motor Cortex Neurons. *Neuron* 93, 929–939.e6. <https://doi.org/10.1016/j.neuron.2017.01.023>
- Roberts, P., Zadan, M., Majidi, C., 2021. Soft Tactile Sensing Skins for Robotics. *Curr. Robot. Rep.* 2, 343–354. <https://doi.org/10.1007/s43154-021-00065-2>
- Rohmer, E., Singh, S.P.N., Freese, M., 2013. V-REP: A versatile and scalable robot simulation framework, in: 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems. Presented at the 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2013), IEEE, Tokyo, pp. 1321–1326. <https://doi.org/10.1109/IROS.2013.6696520>

Sachidhanandam, S., Sreenivasan, V., Kyriakatos, A., Kremer, Y., Petersen, C.C.H., 2013. Membrane potential correlates of sensory perception in mouse barrel cortex. *Nat. Neurosci.* 16, 1671–1677.

<https://doi.org/10.1038/nn.3532>

Shen, Y., Campbell, R.E., Côté, D.C., Paquet, M.-E., 2020. Challenges for Therapeutic Applications of Opsin-Based Optogenetic Tools in Humans. *Front. Neural Circuits* 14, 41.

<https://doi.org/10.3389/fncir.2020.00041>

Weiss, J.M., Flesher, S.N., Franklin, R., Collinger, J.L., Gaunt, R.A., 2019. Artifact-free recordings in human bidirectional brain–computer interfaces. *J. Neural Eng.* 16, 016002. <https://doi.org/10.1088/1741-2552/aae748>

Wodlinger, B., Downey, J.E., Tyler-Kabara, E.C., Schwartz, A.B., Boninger, M.L., Collinger, J.L., 2015. Ten-dimensional anthropomorphic arm control in a human brain–machine interface: difficulties, solutions, and limitations. *J. Neural Eng.* 12, 016011. <https://doi.org/10.1088/1741-2560/12/1/016011>