

Reviewed Preprint

v1 • January 6, 2025

Not revised

Reviewed Preprint

v2 • January 14, 2026

Revised by authors

For correspondence:

thmurphy@mail.ubc.ca

Competing interests:

No competing interests declared

Funding:

See [page 40](#)

Reviewing editor:

Juan Alvaro Gallego, Imperial College London, United Kingdom

© 2025, Gupta & Murphy. This article

is distributed under the terms of the

[Creative Commons Attribution](#)

[License](#), which permits unrestricted

use and redistribution provided that

the original author and source are

credited.

Real-Time Closed-Loop Feedback System For Mouse Mesoscale Cortical Signal And Movement Control: CLoPy

Pankaj K Gupta^{1,2}, Timothy H Murphy^{1,2} 

¹Department of Psychiatry, University of British Columbia, Vancouver, Canada • ²Djavad Mowafaghian Centre for Brain Health (DMCBH), University of British Columbia, Vancouver, Canada

eLife Assessment

This study presents a platform to implement closed-loop experiments in mice based on auditory feedback. The authors provide **convincing** evidence that their platform enables a variety of closed-loop experiments using neural or movement signals, indicating that it will be a **valuable** resource to the neuroscience community. The paper could be strengthened by the addition of additional tutorials, such as on how to run an experiment.

<https://doi.org/10.7554/eLife.105070.2.sa4>

Abstract

Increasingly, experiments designed to provide practical perturbations to circuits or behavior are required for hypothesis testing in various disciplines ranging from motor learning to recovery after injury. We present the implementation and efficacy of an open-source closed-loop neurofeedback (CLNF) and closed-loop movement feedback (CLMF) system. In CLNF, we measure mm-scale cortical mesoscale activity with GCaMP6s and provide graded auditory feedback (within 63 ms) based on changes in dorsal-cortical activation within regions of interest (ROI) and with a specified rule. Single or dual ROIs (ROI1, ROI2) on the dorsal cortical map were selected as targets. Both motor and sensory regions supported closed-loop training in male and female mice. Mice modulated activity in rule-specific target cortical ROIs to get increasing rewards over days (RM ANOVA $p=2.83e-5$) and adapted to changes in ROI rules (RM ANOVA $p=8.3e-10$, Table 4 for different rule changes). In CLMF, feedback (within 67 ms) was based on tracking a specified body movement, and rewards were generated when the behavior reached a threshold. For movement training, the group that received graded auditory feedback performed significantly better (RM-ANOVA $p=9.6e-7$) than a control group (RM-ANOVA $p=0.49$) within four training days. Additionally, mice can learn a change in task rule from left forelimb to right forelimb within a day, after a brief performance drop on day 5. Offline analysis of neural data and behavioral tracking revealed changes in the overall distribution of Ca²⁺ fluorescence values in CLNF and body-part speed values in CLMF experiments. Increased CLMF performance was accompanied by a decrease in task latency and cortical $\Delta F/F_0$ amplitude during the task, indicating lower cortical activation as the task gets more familiar.

Introduction

In neuroscience, brain activity and behavior are often studied as separate entities that do not interact in real time. Assessments are typically performed post hoc, and experimental designs rarely incorporate dynamic feedback based on ongoing neural activity. In contrast, closed-loop systems provide a framework where real-time interaction occurs between the brain, subject, and goal-directed outcomes, enabling precise modulation of neural and behavioral dynamics (Clancy et al. 2014 ). First introduced in the pioneering work of (Fetz 1969 ), closed-loop feedback has since demonstrated significant potential in understanding and modulating brain-behavior

relationships. There has been recent interest in its application to rodent models (Clancy and Mrsic-Flogel 2021 [\[1\]](#); Clancy et al. 2014 [\[2\]](#); Prsa, Galiñanes, and Huber 2017 [\[3\]](#); Srinivasan et al. 2018 [\[4\]](#); Knudsen and Wallis 2020 [\[5\]](#); Paz et al. 2013 [\[6\]](#); Ching et al. 2013 [\[7\]](#); Rosin et al. 2011 [\[8\]](#); Widge and Moritz 2014 [\[9\]](#); Sun et al. 2022 [\[10\]](#)). Some related studies have demonstrated the feasibility of closed-loop feedback in rodents, including hippocampal electrical feedback to disrupt memory consolidation (Girardeau et al. 2009 [\[11\]](#)), optogenetic perturbations of somatosensory circuits during behavior (O'Connor et al. 2013 [\[12\]](#)), and more recent advances employing targeted optogenetic interventions to guide behavior (Abbasi et al. 2023 [\[13\]](#)). Opportunities also exist in extending real time pose classification (Forys et al. 2020 [\[14\]](#); Kane et al. 2020 [\[15\]](#)) and movement perturbation (M. W. Mathis, Mathis, and Uchida 2017 [\[16\]](#)) to shape aspects of an animal's motor repertoire. This gap is due, in part, to technical challenges such as the need for low-latency data acquisition and processing, compact and portable hardware, and seamless integration with animal behavior. Moreover, the inherent complexity of rodent neurophysiology and behavior (e.g., undesired movements, variability in neural signals, and behavioral metrics) further complicates the implementation of closed-loop paradigms.

Closed-loop systems are particularly promising for addressing critical questions in systems neuroscience, such as how brain activity dynamically adapts to feedback or how neural circuits reorganize during motor learning and recovery. Recent advancements, such as genetically encoded calcium indicators and mesoscopic imaging, have opened new avenues for studying these phenomena (Clancy and Mrsic-Flogel 2021 [\[1\]](#); Clancy et al. 2014 [\[2\]](#); Prsa, Galiñanes, and Huber 2017 [\[3\]](#)). However, many of the available closed-loop systems remain technically complex, expensive, or inaccessible to labs with limited resources. Additionally, most existing platforms focus on narrowly defined use cases, limiting their adaptability to diverse experimental paradigms. Nevertheless, recent work has demonstrated behavioral control through feedback, such as in (M. W. Mathis, Mathis, and Uchida 2017 [\[16\]](#)), where mice learned to pull a joystick to a target and adapt this motion under force perturbations.

To address these gaps, we present CLoPy - a flexible, cost-effective, and open-source Python-based platform for closed-loop feedback experiments in neuroscience. CLoPy integrates real-time data acquisition, processing, and feedback delivery, providing researchers with a robust tool to study dynamic brain-behavior interactions in head-fixed mice. The platform supports two modes of graded feedback: (1) Closed-Loop Neurofeedback (CLNF), which derives feedback from neuronal activity, and (2) Closed-Loop Movement Feedback (CLMF), which bases feedback on observed body movements. By combining genetically encoded calcium imaging with behavioral tracking, CLoPy enables researchers to explore questions such as how cortical dynamics contribute to motor learning or how feedback modulates neural plasticity in neurological models.

Crucially, CLoPy is designed with accessibility and versatility in mind. The platform utilizes widely available hardware, including Raspberry Pi and Nvidia Jetson, and leverages Python-based open-source software libraries. These features make CLoPy an ideal resource for the broader open-source neuroscience community, particularly for labs seeking to implement closed-loop systems without substantial financial or technical barriers. Furthermore, CLoPy's modular design allows researchers to adapt it for a wide range of experimental paradigms, from studying cortical reorganization during motor learning to investigating sensory feedback mechanisms. By framing CLoPy within the context of accessible, open-source neuroscience tools, we aim to broaden its impact and utility across the systems neuroscience community.

In this study, we demonstrate the implementation and efficacy of CLoPy in head-fixed mice (Figure 1 [\[17\]](#)), showcasing its ability to deliver precise, low-latency feedback. We provide all software, hardware schematics, and protocols necessary for its adaptation to various experimental scenarios, positioning CLoPy as a valuable tool for advancing closed-loop neuroscience research.




Figure 1. Setup of real-time feedback for GCaMP6 cortical activity and movements

A) GCaMP-based Closed-loop Feedback and Reward System: Mice expressing GCaMP6, with surgically implanted transcranial windows and a head-bar, were positioned beneath the imaging camera, with GCaMP6 excitation light at 470 nm. A secondary wavelength of 440 nm was used for continuous reflectance signals to measure hemodynamic changes, which were then applied to correct the fluorescence signal. The RGB camera was equipped with bandpass filters that allowed only 520 nm epifluorescence and 440 nm reflectance to be simultaneously collected. i) Mice with transcranial windows were head-fixed beneath an imaging camera, with the cortical window illuminated using 440 nm (for reflectance signal used for hemodynamic correction) and 470 nm (for GCaMP excitation) light. Epifluorescence at 520 nm and reflectance at 440 nm were captured at 15 fps using a bandpass filter, integrated within the cortical imaging system. ii) The captured images were simultaneously saved and processed to compute $\Delta F/F_0$ in real-time using a Raspberry Pi 4B model. Pre-selected regions of interest (ROIs) were continuously monitored, and rule-specific activation was calculated based on the $\Delta F/F_0$ signal. The left panel displayed wide-field cortical GCaMP6 fluorescence (green) and reflectance (blue), while the right panel showed the real-time calculated and corrected $\Delta F/F_0$ map, generated using a moving average of the captured images. The target ROIs were marked as green (R1) and red (R2), although a single ROI could also be selected for monitoring. iii) For example, as shown on the $\Delta F/F_0$ map, ROIs R1 and R2 were continuously monitored, and the average activity across these ROIs was calculated. iv) When the task rule was defined as “ $R1 - R2 > \text{threshold}$ ”, the difference between R1 and R2 activities was mapped to a non-linear function that generated graded audio tone frequencies (ranging from 1 kHz to 22 kHz), as illustrated in the figure. Task rules could be modified within the setup on any given day, and the corresponding activation levels were automatically mapped to the audio frequencies. The “threshold”, expressed in $\Delta F/F_0$ units as a percent change, was adjustable based on the experimental design. v) Upon reaching the rule-specific threshold for activity, in addition to the increase in audio tone frequency, a water reward was delivered to the head-fixed mouse. B) Closed-loop Behavior Feedback and Reward Setup: A specialized transparent head-fixation chamber was custom-designed using 3mm-thick plexiglass material (3D model available on GitHub link below) to enable multi-view behavioral imaging and real-time tracking of body parts. The rectangular chamber was equipped with two strategically positioned mirrors—one at the bottom, angled at 45 degrees, and one at the front, angled at 30 degrees—facilitating multi-view imaging of the head-fixed mouse with a single camera. i) A Dalsa CCD camera was connected to a PC for widefield cortical imaging during the session. An additional camera connected to a PC with a GPU was used to capture the behavior video stream of the headfixed mouse. ii) Video frames of the behavior were processed in real-time on a GPU (Nvidia Jetson Orin), which tracked the body parts using a custom trained DeepLabCut-Live (Kane et al., 2020) model. iii) Auditory feedback was provided using a non-linear function mapping paw speeds to corresponding audio tone frequencies. iv) The head-fixed mouse, positioned in the transparent chamber, was able to freely move its body parts, while its behavior was continuously recorded. This setup allowed for the capture of three distinct views of the mouse—side, front, and bottom profiles—and enabled the real-time tracking of multiple body parts, including the snout, left and right forelimbs, left and right hindlimbs, and the base of the tail.

Results

Consistent with prior studies demonstrating that mice can exert volitional control over their brain activity through feedback-based paradigms (Clancy et al. 2014; Luo et al. 2024; Neely, Koralek, et al. 2018; Clancy and Mrsic-Flogel 2021), our experiments show that mice learned to robustly modulate cortical signals when provided with closed-loop neurofeedback. Across training sessions, animals not only acquired the task (Figure 4A, initial rule RM ANOVA $p=2.83e-5$; Figure 4B, No-rule-change group RM-ANOVA $p=9.6e-7$) but also exhibited adaptability to altered feedback rules (Figure 4A, performance drop after rule change ANOVA, $p=8.7e-9$, new rule RM ANOVA $p=8.3e-10$), highlighting the flexibility of their learning. While feedback-driven brain control is well established, our work extends these findings by including multiple cortical regions, and by demonstrating that feedback based on movement-related signals can also drive learning despite the additional challenge of real-time behavioral tracking.

In our closed-loop experiment, mice were trained to associate a specific body movement with a reward. Over the course of 10 training sessions, the mice progressively learned the task, as evidenced by increased task performance across sessions (Figure 4A). The acquisition of this learned behavior was consistent with previous findings on motor learning, where reinforcement and task repetition lead to the refinement of motor skills (Wolpert, Diedrichsen, and Flanagan 2011; Krakauer et al. 2019). Importantly, the reward was provided 1 s after a successful trial to study the timing and nature of cortical activation in response to both the task and the reward. Concurrent with behavioral training, widefield cortical activity was recorded observing distinct patterns of neural activation that evolved as the mice learned the task.

We implemented and tested a compact, cross-platform system designed for closed-loop feedback experiments. Using this system, head-fixed mice successfully learned to associate either cortical activity or movement-derived signals with auditory feedback and water rewards. Performance increased steadily over sessions: in both one-ROI and two-ROI paradigms, animals showed a significant rise in the number of rewards obtained compared to their baseline (day 0) and control groups without feedback (Animations 1A–F). These results indicate that the system effectively links cortical dynamics or movement signals to behaviorally meaningful outcomes. Notably, both neurofeedback (CLNF) and movement-feedback (CLMF) modes produced reliable and graded closed-loop feedback.

In CLNF experiments, auditory tones mapped to $\Delta F/F_0$ values from cortical regions of interest accurately reflected ongoing neural activity (Figure 1). In two-ROI paradigms, differences between ROI activity were similarly mapped to sound frequency, and the commanded audio outputs were verified to be accurate across the full frequency range (Supplementary Figure 2). Improvements to the audio generation pipeline eliminated previous nonlinearities above 10 kHz, confirming that the feedback signal matched the intended output. Importantly, feedback latency was tested and found to be well within the range necessary for effective closed-loop learning. Using LED-triggering experiments (Supplementary Figure 1), we confirmed that the Raspberry Pi-based CLNF system delivered graded auditory feedback with an average latency of 63 ± 15 ms from event detection to output, using GPIO pins. Please note this latency calculation used GPIO pins for delay measurements, and there could be additional minimal overheads to generate the actual audio feedback signals. These findings demonstrate that our system provides rapid and accurate feedback sufficient to drive volitional control of brain or movement signals in head-fixed mice.

For the CLMF experiments, an Omron Sentech STC-MCCM401U3V USB3 Vision camera, connected to the Nvidia Jetson via its Python software developer kit (SDK), was used to calculate the feedback delay. The incoming stream of frames was processed in real-time using a custom deep-neural-network model that was trained using DeepLabCut (A. Mathis et al. 2018) and DeepLabCut-Live (Kane et al. 2020), designed to track previously defined points on the mouse. The model was incrementally improved by fine-tuning and re-training 26 times, using 2080 manually labeled frames spanning 52 videos of 10 mice. The pre-trained model, along with the data, is available at the link shared in the data and code availability statement in methods for anyone to use and fine-

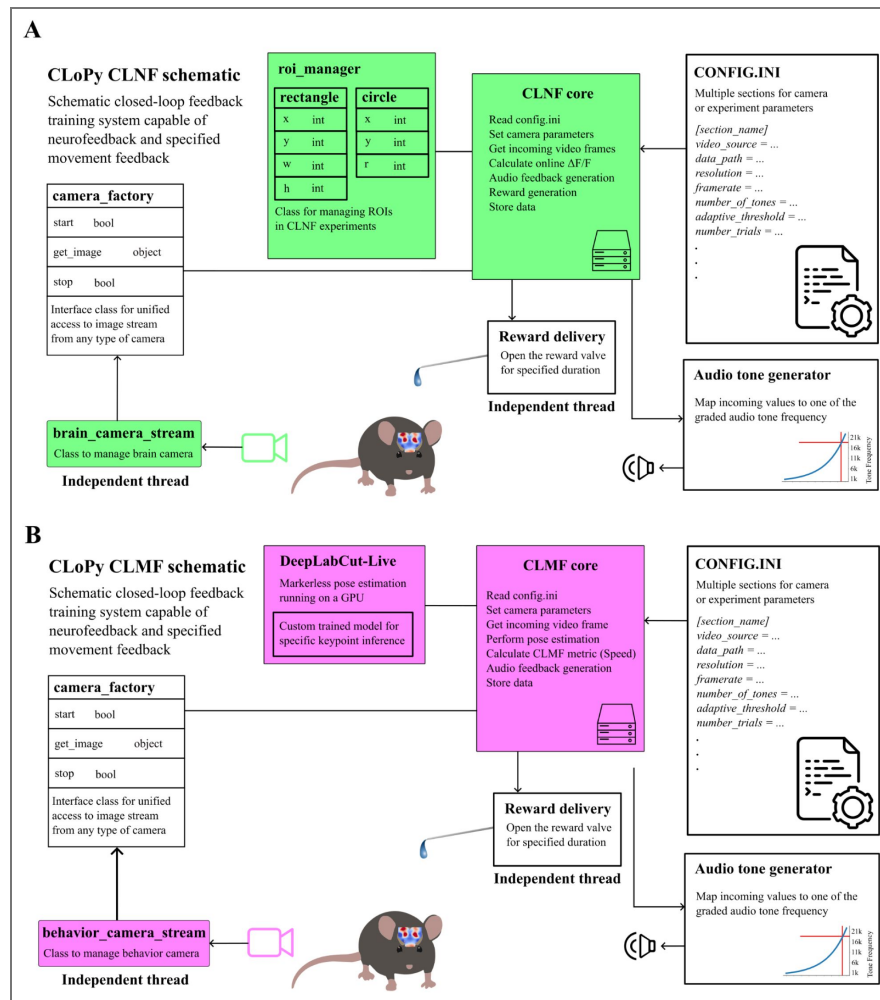


Figure 2. Schematic of the closed-loop feedback training system (CLOPy) for neurofeedback and specified movement feedback.

A) The components of CLOPy CLNF system are presented in a block diagram. Modular components such as the configuration file, camera factory, audio tone generator, and reward delivery system are displayed and are utilized by both the Closed-Loop Neurofeedback (CLNF) and Closed-Loop Movement Feedback (CLMF, in B) systems. The configuration file (config.ini) stored all configurable parameters of the system, including camera settings, feedback parameters, reward thresholds, number of trials, and the duration of trial and rest periods, under an experiment-specific section. The camera factory was an abstract class that provided a unified interface for a programmable camera to the CLNF and CLMF(in B) systems. This abstraction allowed the core system to remain independent of the specific camera libraries required for image streaming. Camera-specific routines were implemented in separate “brain_camera_stream” and “behavior_camera_stream” classes, which inherited functions from the “camera_factory” superclass and ran in independent thread processes. The Region of Interest (ROI) manager was used by the CLNF core to maintain a list of ROIs, as well as routines to perform rule-specific operations on them, as specified in config.ini. An ROI could be defined as a rectangle (with the upper-left corner coordinates, height, and width) or as a circle (with center coordinates and radius). The audio tone generator mapped the target activity (fluorescence signal in CLNF) to graded audio tone frequencies. It generated audio signals at 44.1 kHz sampling based on the specified frequency and sent the signal to the audio output. Reward delivery was controlled by opening a solenoid valve for a specified duration, which was managed in a separate process thread. The CLNF core was the main program responsible for running the CLNF system. It utilized config.ini, the camera factory, the ROI manager, and integrated the audio tone generator and reward delivery functions. The system also saved the recorded data and configuration parameters with unique identifiers. B) The components of CLOPy CLMF system are presented in a block diagram. The common components of the setup are described in A. The audio tone generator mapped the target activity (control-point speed in CLMF) to graded audio tone frequencies. It generated audio signals at 44.1 kHz sampling based on the specified frequency and sent the signal to the audio output. The CLMF core was the primary program responsible for operating the CLMF system. It utilized config.ini, the camera factory, and DeepLabCut-Live(Kane et al. 2020), integrating the audio tone generator and reward delivery functions. This module also saved the data and configuration parameters with unique identifiers.

tune to adapt for similar platforms. The model was integrated into our CLMF program and deployed on an Nvidia Jetson device for real-time inference of tracked points. A Python-based virtual environment using Conda was created to install all software dependencies. The coordinates of the tracked points for each frame were appended to a list, forming a temporal sequence referred to as “tracks.” Upon offline analysis, these real-time tracks were found to be both accurate and stable throughout the duration of the videos.

To calculate the feedback delay for movements, a red LED was placed within the behavior camera’s field of view. Whenever a threshold-crossing movement was detected in the real-time tracking, the system triggered the LED to turn on. Temporal traces of the tracked left forelimb and the pixel brightness of the red LED were then extracted from the recorded video. By comparing these traces, the average delay between the detected movements and the LED illumination was calculated to be 67 ± 20 ms, which represents the delay in the closed-loop feedback for the CLMF experiments.

For CLMF, mice were headfixed in a specially designed head-fixing chamber (Figure 1Biii [↗](#), design guide and 3D model in methods) to achieve multiview behavioral recording from a single camera using mirrors (see methods section). In brief, mice are headfixed in a transparent rectangular tunnel (top view) with a mirror at the front (front view) and at the bottom (bottom view) that allows multiple views of the mouse body. The body parts that we track are: snout-top (snout in the top view), tail-top (base of the tail in the top view), snout-front (snout in the front view), FLL-front (left forelimb in the front view), FLR-front (right forelimb in the front view), snout-bottom (snout in the bottom view), FLL-bottom (left forelimb in the bottom view), FLR-bottom (right forelimb in the bottom view), HLL-bottom (left hindlimb in the bottom view), HLR-bottom (right hindlimb in the bottom view), tail-bottom (base of the tail in the bottom view) (Figure 1B [↗](#)). The rationale for selecting these body parts in a particular view was that they needed to be visible at all times to avoid misclassification in real-time tracking. By combining the tracks of a body part in different views, we can form a 3D track of the body part. In a 3D coordinate space having X, Y, and Z axes, tracked points (xb, yb) in the bottom view were treated as being in the 3D X-Y plane, and tracked points (xf, yf) in the front view were treated as being in the 3D Y-Z plane. Thus, $X = xb$, $Y = yb$, $Z = yf$ formed tracked points in 3D for a given body part that was tracked in multiple views.

For example, FLL-front and FLL-bottom were tracking the left forelimb in front and bottom views, and by combining the tracks of these two points, we obtained a 3D track of the left forelimb (Figure 3E [↗](#)). Although these 3D tracks are available in real-time, for our CLMF experiments, we used 2D tracks (xb, yb) for behavioral feedback to keep the task simpler and comparable across mice.

Audio output channels and GPIO pins on the Nvidia Jetson were used for audio feedback and reward delivery, respectively. Tracks of each body part were used to calculate the speed of those points in real-time, and a selected body part (also referred to as a control point) was mapped to a function generating proportional audio frequency (same as in CLNF, details in the method section). In the software we have developed, one can also choose to calculate acceleration, radial distance, angular velocity, etc., from these tracks and map it to the function generating varying audio frequency feedback. For our work, a range of spontaneous speeds was calculated from a baseline recording before the start of the training. The threshold speed to receive a reward was also calculated from the baseline recording and set at a value that would have yielded a reward rate of one reward per two-minute period (translates to a basal performance of 25% in our trial task structure). This thresholding step was done to allow the mice to discover task rules and keep them motivated.

Both CLNF and CLMF experiments shared a similar experimental protocol (surgery, habituation, then several days of training, Figure 3A [↗](#)). A daily session starts with a 30 sec rest period (no rewards or task feedback) followed by 60 trials (maximum 30 sec each) with a 10 sec rest between trials in CLNF, and a minimum of 10 sec rest or until tracked points are stable for 5 sec in CLMF (Figure 3B [↗](#), 3D [↗](#)). After the habituation period, a spontaneous session of 30 minutes was recorded where mice were not given any feedback or rewards.




Figure 3. Experimental protocol and trial structure

A) Experimental Protocol (detailed in Methods): In brief, 90-day-old transgenic male and female mice were implanted with a transcranial window and allowed to recover for a minimum of 7 days. One day before the start of the experiment, the mice were placed on a water-restriction protocol (as detailed in Methods). Closed-loop experiment training commenced on day 1, during which mice were required to modulate either their target brain activity (GCaMP signals in regions of interest) or target behavior (tracked paw-speed) during daily sessions of approximately 45 minutes over the course of 10 days. Throughout this period, both cortical and behavioral activities were recorded. After the final experimental session, the mice were removed from the water-restriction protocol. B) Trial Structure of Cortical GCaMP-based Feedback Sessions: Each trial was preceded by a minimum of 10 seconds of rest, which was extended if the tracked body parts of the mouse were not stable (sum of changes across body parts > 1.5 mm i.e. 5 pixel values). Once the mouse remained stable and refrained from moving its limbs, the trial began with a basal audio tone of 1 kHz. The mice then had 30 seconds to increase rule-based activations (in the selected ROI) up to a threshold value to receive a water reward. A trial ended as soon as the activation reached the threshold, triggering a reward delivery for success, or timed out after 30 seconds with an overhead buzzer serving as a negative signal of failure. Both the reward and the negative signal were delivered within 1 second after the audio ceased at the end of each trial. C) Example dorsal cortical $\Delta F/F_0$ activity was recorded and overlaid with a subset of Allen CCF coordinates, which could be selected as the center of candidate regions of interest (ROIs) shown as Green and Pink color squares. D) Trial Structure of Behavior Feedback Sessions: The behavioral feedback trials followed a similar structure to the cortical feedback trials, with each trial preceded by at least 10 seconds of rest. The trial began with a basal tone of 1 kHz, which increased in frequency as the mouse's paw speed increased. E) Example forelimb Tracking during Feedback Sessions: Forelimb tracking was performed in both the left (blue) and right (green) forelimbs using a camera coordinate system on day 4 of training with mouse FM2, which received feedback based on left forelimb speed. The forelimbs were tracked in 3D, leveraging multiple camera views captured using mirrors positioned at the bottom and front of the setup (Figure 1B).




Figure 4. Closed-loop feedback helped mice learn the task and achieve superior performance in CLNF and CLMF in both experiments.

A) No-rule-change (N=23, n=60, in Table 5) and Rule-change (N=17, n=60) mice were able to learn the CLNF task over several sessions, with performance above 70% by the 10th session (RM ANOVA $p=2.83e-5$). The rule change (in pink, day 11) led to a sharp decline in performance (ANOVA, $p=8.7e-9$), but the mice were able to adapt and learn the task rule change (RM ANOVA $p=8.3e-10$; see Table 5 for different rule changes). The method to determine the ROI(s) used in the changed task rule is described in the methods section. B) Three groups were employed for CLMF experiments. The “Rule-change” group (N=8, n=60 received feedback, in pink) was trained with task rule mapping auditory feedback to the speed of the left forelimb and was able to perform above a 70% success rate in four days. The task rule mapping was changed from the left to the right forelimb on day 5, so the rewards as well as audio frequencies would now be controlled by the right forelimb. The “No-rule-change” group (N=4, n=60 received audio feedback, no rule change, in green) and the “No-feedback” group (N=4, n=60 no graded audio feedback, no rule change, in blue) were control groups to investigate the role of audio feedback. The performance of the “No-feedback” mice, who did not receive the graded feedback, was never on par (RM-ANOVA $p=0.49$) with the “No-rule-change” group that received the feedback (RM-ANOVA $p=9.6e-7$). Additionally, to test the effect of auditory feedback on mice that already had learned the task, we turned off the auditory feedback on day 10 for all mice (indicated by the speaker with a cross). There was no significant change in performance due to feedback removal, indicating that feedback was not necessary once the task was learned. C) Task latencies in each group follow the trend of their performance. Rule change (N=8) and no rule change (N=4) task latencies gradually came down, with an exception on day 5 for rule change when the task rule was changed. No feedback (N=4) task latencies are never on par with the groups that received feedback. D) CLMF Rule-change (N=8) behavior, we looked at the maximum speeds of the left and right forelimbs. The paw with the maximum speed follows the task rule and switches with the change in the task rule. It is worth noting that the task was not restricted to other body parts; i.e., the mice were free to move other body parts along with the control point.

The spontaneous session was used to establish baseline levels of GCaMP activity (in target ROI(s) for CLNF experiments) or speed of a target body part for CLMF experiments. This was done to calculate the animal-specific threshold for future training sessions. A success in the trial resulted in a water drop reward that was delivered 1 sec after the end of the trial, and a failed trial ended with a buzzer vibrator 1 sec after the end of the trial (Figure 3B, 3D).

Mice can explore and learn an arbitrary task, rule, and target conditions

CLNF training (Figure 1A) required real-time image processing, feedback, and reward generation. Feedback was a graded auditory tone mapped to the relative changes in selected cortical ROI(s) or movement speed of a tracked body part. Training was conducted using multiple sets of cortical ROI(s) (list of ROI rules and their description can be found in Table 3 and visualized in Supplementary Figure 7) on both male and female mice (see Table 5 for list of mice and their task rules), wherein the task was to increase the activity in the selected ROI(s), according to the rule (also referred to as ‘the rule’ in the future), to exceed a predetermined threshold (Figure 6A, 6B). Fluorescence activity changes ($\Delta F/F_0$) were calculated using a running baseline of 5 seconds, with single or dual ROIs on the dorsal cortical map selected as targets. The distribution of target ROI(s) activity on day 9 of training was higher than on day 1 (Figure 5A). A linear regression analysis of whole-session $\Delta F/F_0$ activity in the 2-ROI experiment revealed a shift in the regression line slope, favoring the ROI that was required to increase according to the task rule (Figure 6C). In general, all ROIs assessed (Table 3, Supplementary Figure 7) that encompassed sensory, pre-motor, and motor areas were capable of supporting increased reward rates over time based on analysis of group and pooled ROI data (Figure 4A, Animation 1). A $\Delta F/F_0$ threshold value was calculated from a baseline session on day 0 that would have allowed 25% performance. Starting from this basal performance of around 25% on day 1, mice (CLNF No-rule-change, N=23, n=60 and CLNF Rule-change, N=17, n=60) were able to discover the task rule and perform above 80% over ten days of training (Figure 4A, RM ANOVA $p=2.83e-5$), and Rule-change mice even learned a change in ROIs or rule reversal (Figure 4A, RM ANOVA $p=8.3e-10$, Table 5 for different rule changes). There were no significant differences between male and female mice (Supplementary Figure 3A). To investigate whether certain ROI(s) were better than others in terms of performance, we performed linear regression of the success rate over the days and, based on the slope of the fitted line, discovered ROI rules that yielded statistically different progression (fast and relatively slower) of success rate from the mean slope of all ROIs (Supplementary Figure 3C). We visualized these significantly different ROI-rule-based success rates and segregated them into fast and slow based on mean slope (≥ 0.095 slope was designated fast, else slow) of the progression (Supplementary Figure 3D). Our analysis revealed that certain ROI rules (see description in methods) lead to a greater increase in success rate over time than others (Supplementary Figure 3D). As mice learned the CLNF task and their performance improved, task latency progressively decreased (Figure 4C). Here, task latency refers to the time required by the mice to complete the task within the 30-s trial window. For this measure, we included all trials, both successful and unsuccessful. In the 2-ROI experiment where the task rule required, for example, “ROI1 - ROI2” activity to cross a threshold for reward delivery, mice displayed divergent strategies. Some animals predominantly increased ROI1 activity, whereas others reduced ROI2 activity, both approaches could lead to successful threshold crossing (Figure 11; see average reward centered responses).

In CLMF training, the real-time speed of a selected (also referred to as control point (CP) in terms of tracked points) point was mapped to the graded audio tone generator function. We trained the mice with the FLL bottom as the CP for auditory feedback and reward generation (example trial in Animation 2). Mice reached 80% performance (Figure 4B, CLMF Rule-change, No-rule-change) in the task within four days of training (RM ANOVA, $p = 8.03e-7$). They were eliciting the target behavior (i.e., moving CP at high speed) more frequently on later days compared to the first day of training (Figure 5B), and the correlations of CP speed profiles became more pronounced during the trial periods as compared to the rest period (Figure 8A).




Figure 5. Closed-loop feedback helped mice learn the task and achieve superior performance in CLNF and CLMF in both experiments.

A) Reward-aligned average ($N=4$) $\Delta F/F_0$ signals associated with the target rule on day1 and day9 (top plot). Kernel density estimate (KDE) of target $\Delta F/F_0$ values during the whole session on day1 and day9 of 1-ROI experiments (bottom plot). B) Reward-aligned average ($N=4$) target paw speed on day1 and day10 (top plot). Kernel density estimate (KDE) of target paw speeds on day1 and day10 (bottom plot). C) In the context of CLNF 2-ROI experiments, bivariate distribution of ROI1 and ROI2 $\Delta F/F_0$ values during whole sessions on day9 and day19, with densities projected on the marginal axes. The task rule on day9 was “ROI1-ROI2 > thresh.” as opposed to “ROI2-ROI1 > thresh.” on day19. The bivariate distribution is significantly different (Multivariate two-sample permutation test, $p=2.3e-12$) on these days, indicating a robust change in activity within these brain regions. D) Joint (bivariate) distribution of left and right paw speeds during the whole session on day4 and day10 of CLMF. Left and right forelimbs were control-point (CP) on day4 and day10 respectively. There is a visible bias in the bivariate distribution towards the CP on respective days.




Figure 6. CLNF: Cortical activity during the closed-loop-neurofeedback training.

A) Target-rule-based $\Delta F/F_0$ traces in green on day 1 with rule-1 (top row), on day 10 with rule-1 (second row), on day 11 with new rule-2 (third row), and on day 19 with rule-2 (fourth row). Shared regions are trial periods and regions between grey areas are rest periods. The grey horizontal line depicts the threshold above which mice would receive a reward. Golden stars show the rewards received, and short vertical lines in black show the spout licks. B) Representative reward-centered average cortical responses of the 2ROI experiment on labeled days. ROI1 (green) and ROI2 (pink) are overlaid on the brain maps. The task rule on day 1 and day 4 was “ROI1-ROI2 > thresh,” as opposed to “ROI2-ROI1 > thresh” on day 11 and day 19. C) Linear regression on ROI1 and ROI2 $\Delta F/F_0$ during whole sessions for an example mouse. The regression fit inclines towards ROI1 in sessions where the rule was “ROI1-ROI2 > thresh.” (day1 slope=0.44, day6 slope=0.32) while it leans towards ROI2 after the task rule switches to “ROI2-ROI1 > thresh.” (day11 slope=0.76, day17 slope=0.80).

Mice can rapidly adapt to changes in the task rule

We have assessed how mice respond to changes in closed-loop feedback rules. CLNF Rule-change mice went through a change in regulated cortical ROI(s) (Figure 6A [↗](#), 6B [↗](#)) after being trained on an initial set of ROI(s) (Table 5 [↗](#), Rule-change). The new ROI(s) were chosen from a list of candidate ROI(s) for which the cortical activations were still low (see “Determining ROI(s) for change in CLNF task rule” under Methods). Reward threshold values were then recalculated for the ROI(s). A full list of ROI(s) we tried in separate mice is listed in Table 5 [↗](#). Mouse performance degraded on the first day of the rule switch to the new ROIs (ANOVA, $p=8.7e-9$), compared to the previous day (Figure 4A [↗](#)), and quickly recovered within five days (RM ANOVA $p=8.3e-10$) of further training. All new cortical ROI(s) appeared to support recovery to the pre-perturbation success rate (Table 5 [↗](#)), and data were pooled across mice. This can also be observed in the bivariate distribution of a whole session $\Delta F/F_0$ values in ROI1 vs ROI2 on day 9 (before rule change) and day 19 (last day after rule change) in Figure 5C [↗](#). Additionally, the linear regression analysis of whole-session $\Delta F/F_0$ activity after rule change in the 2-ROI experiment revealed a shift in the regression line slope, favoring the ROI that was required to increase according to the new task rule (Figure 6C [↗](#)). Following the rule change, CLNF task latency increased substantially and persisted at a high level across subsequent sessions, eventually declining as task performance adapted to the new rule (Figure 4C [↗](#)).

Similarly, in the CLMF experiments, the target body part was changed from FLL to FLR for Rule-change mice on day 5 and was associated with a significant drop in their success rate from 75-80% to ~40% (ANOVA $p=0.008$, Figure 4B [↗](#)). Surprisingly, mice quickly adapted to the rule change and started performing above 70% on day 6 (example trial in Animation 3 [↗](#)). Looking closely at their reward rate on day 5 (day of rule change), they had a higher reward rate in the second half of the session as compared to the first half (Supplementary Figure 4A [↗](#)), indicating they were adapting to the rule change within one session. This can also be observed in the bivariate distribution of left and right paw speeds on day 4 (before rule change) and day 10 (last day after rule change) in Figure 5D [↗](#). As the mice were learning the CLMF task (increasing performance), task latency decreased (Figure 4D [↗](#)). Task latency in this context is the time taken by mice to perform the task within a trial period of 30 s. We included all trials, both successful and unsuccessful, in our calculation. Given that the maximum trial duration is 30 s, the task latency was capped at 30 s. Following the trend in task performance, the task latency started decreasing during the first four days but increased on day 5 (rule change) and started to drop again afterward.

We also examined the average paw speeds and distributions during trial periods. It is worthwhile noting that for the rule-change group, left forelimb speeds were higher than right forelimb speeds from day 1 to day 4 (when task rule was to move left forelimb). When the rule was switched from left forelimb to right forelimb on day 5, left forelimb speeds dropped below the right forelimb speeds (Supplementary Figure 4E [↗](#)).

Graded feedback helps in task exploration during learning, but not after learning

To investigate the role of audio feedback in our task, we also had a group with no task-related graded audio feedback (CLMF No-feedback) and instead received audio with a constant frequency (1 kHz) throughout the trials. CLMF No-rule-change mice who received continuous graded auditory feedback significantly improved their task performance (CLMF No-rule-change RM-ANOVA $p=9.6e-7$) and outperformed the CLMF No-feedback mice (Figure 4B [↗](#)) very early (No-feedback RM-ANOVA $p=0.49$), indicating the positive role of graded feedback for task exploration and learning the association. When graded audio feedback was removed for CLMF Rule-change mice on day 10, it did not affect their task performance, indicating the feedback was not essential to keep performing the task that they had already learned.

Cortical responses became focal and more correlated as mice learned the rewarded behavior in CLMF

There are reports of cortical plasticity during motor learning tasks, both at cellular and mesoscopic scales (Makino et al. 2017 [↗](#); Huber et al. 2012 [↗](#); W. E. Allen et al. 2017 [↗](#)). As mice become proficient in a task, their brain activity becomes more focused and less widespread (Figure 9 [↗](#)). We noticed the peak activity in different regions of the cortex around the rewarding movement decreased $\Delta F/F_0$ in amplitude over days (Figure 9A [↗](#), 9B, 9C, 9E). To quantify this, we measured the peak $\Delta F/F_0$ ($\Delta F/F_0$ peak) value in the time window from -1s to +1s relative to the body part speed threshold crossing event in each cortical region. Consistent with our observations, $\Delta F/F_0$ peak in several cortical regions gradually decreased over sessions, including the olfactory bulb, sensory forelimb, and primary visual cortex (Figure 9A [↗](#), 9B, 9C). Notably, when the task rule was changed from FLL to FLR in the CLMF Rule-change mice, we observed a significant increase in $\Delta F/F_0$ peak in regions including the olfactory bulb (OB), forelimb cortex (FL), hindlimb cortex (HL), and primary visual cortex (V1), as shown in Figure 9 [↗](#). We believe the decrease in $\Delta F/F_0$ peak is unlikely to be driven by changes in movement, as movement amplitudes did not decrease significantly during these periods (Figure 9D [↗](#) CLMF Rule-change). However, the decrease in $\Delta F/F_0$ peak followed the same trend as task latency (Figure 4D [↗](#)), suggesting that the decrease in $\Delta F/F_0$ peak is especially prominent in trials where mice were prepared to make the fast movement.

These results suggest that motor learning led to less cortical activation across multiple regions, which may reflect more efficient processing of movement-related activity. In addition to $\Delta F/F_0$ peak reflecting signs of potentially more efficient cortical signaling, intracortical GCAMP transients measured over days became more correlated (to each other) during the trials as the task performance increased over the sessions (Figure 9E [↗](#)). Analysis of pairwise correlations between cortical regions (referred to as seed pixel correlation maps) revealed distinct network activations during rest and trial periods (Figure 10 [↗](#)). While the general structure of the correlation maps remained consistent between trial and rest periods, certain correlations, such as between V1 and RS, were heightened during task trials, whereas others, such as between M1, M2, FL, and HL, consistently increased over the training sessions (Figure 10 [↗](#)).

To statistically examine the differences in seed pixel correlations during trials and rest periods, as well as how these correlations changed over training sessions (Day 1-10), we conducted a two-way repeated measures ANOVA (RM-ANOVA) on the seed pixel correlation maps for each day. The two variables for the RM-ANOVA were experimental condition (trial vs. rest) and session number (Day 1-10). This analysis generated two distinct matrices of Bonferroni-corrected p-values (Figure 10 [↗](#)), one corresponding to each variable, which segregated the seed pixel correlations that were different between trial and rest periods and those that changed over the training sessions.

Distinct task- and reward-related cortical dynamics

During the early sessions (days 1 to 3), cortical activity was observed to be spatially widespread and engaged multiple cortical regions. Temporally, the activity spanned both task-related and reward-related events, with no clear distinction between the two phases (Figure 9B [↗](#), left). This broad activation pattern is consistent with the initial stages of learning, where the brain recruits extensive cortical networks to process novel tasks and integrate sensory feedback (Peters, Chen, and Komiyama 2014 [↗](#)).

As the mouse performance improved in the later sessions (Days 8 to 10), the cortical activity became more segregated both spatially and temporally (Figure 9B [↗](#), middle). This segregation was particularly notable in mice that received closed-loop feedback (Rule-change), where the spatiotemporal patterns of cortical activation were more closely aligned with the specific task and reward events. This transition from widespread to segregated activation is indicative of the brain's optimization of neural resources as the task becomes more familiar, a phenomenon that has been reported in studies of skill acquisition and motor learning (Makino et al. 2017 [↗](#); Huber et al. 2012 [↗](#)). Previous studies have shown that feedback, especially when provided in a temporally





Figure 7. CLMF: Speed of the tracked target body part and cortical activity.

A) Left forelimb speed (black), target threshold (grey line), and rewards (golden stars) during a sample period in a session on day 1 of the closed-loop training (top row). Shaded regions are trial periods with interspersed rest periods in white. Left forelimb speed, and rewards on day4 (second row). The target body part was changed from the left forelimb to the right forelimb on day5 (third row). Thus, day5 is the first training day with the new rule. Right forelimb speed, and rewards on day 10 of the training (fourth row). B) Reward-centered average cortical responses on days corresponding to rows in A. The target threshold was crossed at -1 s, and the reward was delivered at 0 s. Notice the task rule change on day 5.



A) Top row: During rewarded trials, over the training sessions of CLMF Rule-change (left), No-rule-change (middle), and No-feedback (right). High correlations (dark cells) between speed profiles of CP indicate a unilateral bias in the movement. It is worth noting the drastic changes in correlations as the control-point was changed from left forelimb to right forelimb in Rule-change mice on day4. Bottom row: During rest periods, over the training sessions of CLMF Rule-change (left), No-rule-change (middle), and No-feedback (right).




Figure 9. Cortical dynamics and network changes during longitudinal CLMF training.

A) Reward-centered average responses in the olfactory bulb decrease over the days as performance increases. Data shown is a representative example from sessions of mouse. B) Cortical responses become focal and closely aligned to the paw movement (green line) and reward (cyan line) events on day 10 for group-1 (received feedback) as compared to group-3 (no-feedback). C) $\Delta F/F_0$ peak values during successful trials (pink) and during rest (cyan) over the ten-day training period in the olfactory bulb (OB) (left, day 1-day 4 p-value=0.025, day 4-day 5 p-value=0.008), forelimb area (FL) (center, day 1-day 4 p-value=0.008, day 4-day 5 p-value=0.04), and primary visual cortex (right, day 1-day 4 p-value=0.04, day 4-day 5 p-value=0.002). Statistical significance was assessed using the Mann-Whitney test and corrected for multiple comparisons with the Benjamini-Hochberg procedure. D) Average movement (mm) of different tracked body-parts during trials in CLMF Rule-change (left), No-rule-change (center), No-feedback (right). E) Correlations between cortical activation on each training session in barrel-cortex (BC, top left), anterolateral motor cortex (ALM, top right), secondary motor cortex (M2, bottom left), and retrosplenial cortex (RS, bottom right).




Figure 10. CLMF cortex-wide seed pixel correlation maps.

Pairwise correlations between activity at cortical locations (also referred to as seed pixel locations). A) Top row: No-rule-change average seed pixel correlation map during trial periods (left), during rest periods (middle), and difference of average correlation map during trial and rest (right). Bottom row: No-feedback average seed pixel correlation map during trial periods (left), during rest periods (middle), and difference of average correlation map during trial and rest (right). B) Significant increase in pairwise seed pixel correlations as RM-ANOVA p-value (Bonferroni corrected) matrix between training sessions over the days (left) and between trial vs rest periods (right) for CLMF No-rule-change mice. C) Significant increase in pairwise seed pixel correlations as RM-ANOVA p-value (Bonferroni corrected) matrix between training sessions over the days (left) and between trial vs rest periods (right) for CLMF No-feedback mice.

precise manner, can accelerate the cortical plasticity associated with learning (Ganguly and Carmena 2009 [↗](#)). Overall, these findings highlight the importance of closed-loop feedback in motor learning and suggest that real-time neurofeedback can enhance the specificity of cortical representations, potentially leading to more efficient and robust learning outcomes.

Discussion

Flexible, cost-effective, and open-source system for a range of closed-loop experiments

We developed a versatile, cost-efficient, and open-source platform designed to support a wide array of closed-loop neuroscience experiments, including closed-loop neural feedback (CLNF) and closed-loop movement feedback (CLMF). This system enhances accessibility for the broader research community, enabling the execution of complex experimental paradigms with minimal resource constraints. Our system is built on readily available hardware components, such as Raspberry Pi and Nvidia Jetson platforms, and leverages Python-based software for real-time data processing, analysis, and feedback implementation. The modular approach ensures that the system can be easily customized to meet the specific requirements of different experimental paradigms. Our study demonstrates the effectiveness of a versatile and cost-effective closed-loop feedback system for modulating brain activity and behavior in head-fixed mice. By integrating real-time feedback based on cortical GCaMP imaging and behavior tracking, we provide strong evidence that such closed-loop systems can be instrumental in exploring the dynamic interplay between brain activity and behavior. The system's ability to provide graded auditory feedback and rewards in response to specific neural and behavioral events showcases its potential for advancing research in neurofeedback and behavior modulation.

The hardware backbone of our system is designed around the Raspberry Pi 4B+ and Nvidia Jetson devices, chosen for their compactness, low cost, high computational power, and wide community support. These platforms have been demonstrated to be effective in neuroscience research, for several applications (Michelson et al. 2023 [↗](#); Murphy et al. 2020 [↗](#); Silasi et al. 2018 [↗](#); Dhillon et al. 2021 [↗](#)). By utilizing open-source software frameworks such as Python, the system can be modified and extended to incorporate additional functionalities or adapt to new experimental protocols.

Our Python-based software stack includes libraries for real-time data acquisition, signal processing, and feedback control. For instance, we utilize OpenCV for video processing and DeepLabCut-Live (A. Mathis et al. 2018 [↗](#); Kane et al. 2020 [↗](#)) for pose estimation, which are essential components in experiments requiring precise monitoring and feedback based on animal behavior. Additionally, we have integrated libraries for handling neural data streams, such as NumPy and SciPy, which facilitate the implementation of complex experimental designs (Akam et al. 2022 [↗](#)) involving multiple data sources and feedback modalities. By integrating the system with LED drivers and opsin-expressing transgenic mouse lines, it is straightforward to achieve precise temporal control over neural activation, enabling the study of causal relationships between neural circuit activity and behavior (Lee et al. 2020 [↗](#)).

The cost-effectiveness of our system is a significant advantage, making it accessible to a broader range of research labs, including those with limited funding. The use of off-the-shelf components and open-source software drastically reduces the overall cost compared to commercially available systems, which often require expensive proprietary hardware and software licenses (C. Allen and Mehler 2019 [↗](#); White et al. 2019 [↗](#)). The open-source nature of our closed-loop platform fosters collaboration and innovation within the research community. By providing unrestricted access, other labs can freely modify, enhance, and share the system, promoting a culture of transparency and shared progress. Importantly, our goal is to increase the number of studies utilizing this platform, enabling a broader exploration of closed-loop methodologies across diverse research contexts. This approach not only accelerates the pace of discovery but also enhances reproducibility and adaptability, ensuring the platform remains a robust tool for advancing neuroscience research.




Figure 11. Average reward centered activity in ROIs for 2-ROI experiment.

Examples of reward centered activity in both the ROIs in a 2 ROI experiment reveals different strategies adopted by mice. For a task rule specifying “ROI1 - ROI2 > threshold”, mouse BZ1 and BZ2 decreased ROI2 activity to meet the reward threshold while mouse BU2 and BU3 increased ROI1 activity.

The system's flexibility is demonstrated by its successful application across two closed-loop experimental paradigms (CLNF, CLMF). For example, in our study, we utilized the system to implement real-time feedback based on intracellular calcium-induced fluorescence imaging in awake, behaving mice. The system provided auditory feedback in response to changes in cortical activation, allowing us to explore the role of real-time feedback in modulating both neural activity and behavior.

Importance of closed-loop feedback systems

Closed-loop feedback systems have gained recognition for their ability to modulate neural circuits and behavior in real time, an approach that aligns with the principles of motor learning and neuroplasticity. The ability of mice to learn and adapt to tasks based on cortical or behavioral feedback underscores the parallels between closed-loop systems and natural proprioceptive mechanisms, where continuous sensory feedback is crucial for motor coordination and spatial awareness. This study builds upon the foundational work of (Fetz 1969 [↗](#)), who first explored the potential of closed-loop neurofeedback, and expands its application to modern neuroscience by incorporating optical brain-computer interfaces.

Our findings are consistent with previous research showing that rodents can achieve volitional control over externally represented variables linked to their behavior (Clancy et al. 2014 [↗](#); Neely, Piech, et al. 2018; Prsa, Galíñanes, and Huber 2017 [↗](#); Clancy and Mrsic-Flogel 2021 [↗](#)). Moreover, the rapid adaptation observed in mice when task rules were altered demonstrates the system's capacity to facilitate learning and neuroplasticity, even when the conditions for achieving rewards are modified. The quick recovery of task performance after a rule change, as evidenced by the improved performance within days of training, highlights the robustness of the closed-loop feedback mechanism. We have focused on large regional cortical GCAMP signals that are relatively slow in kinetics.

Neuroplasticity and cortical dynamics

The observation that cortical responses became more localized as mice learned the rewarded behavior aligns with established theories of motor learning, where neural efficiency improves with practice (Makino et al. 2016 [↗](#); Huber et al. 2012 [↗](#); C. B. Allen, Celikel, and Feldman 2003 [↗](#)). As mice became more proficient in the task, the widespread cortical activity observed during the initial training sessions became more regionally localized, indicating more efficient neural processing. The reduction in $\Delta F/F_0$ peak values across sessions suggests that the brain becomes more efficient at processing task-relevant information, a phenomenon consistent with the optimization of neural circuits observed in skilled motor learning (Wolpert, Diedrichsen, and Flanagan 2011 [↗](#); Krakauer et al. 2019 [↗](#)).

The distinct spatiotemporal patterns of cortical activation observed in mice receiving closed-loop feedback further support the role of real-time feedback in enhancing cortical plasticity. The pronounced segregation of task-related and reward-related cortical dynamics in the later training sessions indicates that closed-loop feedback facilitates the refinement of neural circuits involved in motor learning. These findings are in line with previous studies that have demonstrated the importance of temporally precise feedback in accelerating cortical reorganization and enhancing learning outcomes (Ganguly and Carmena 2009 [↗](#)).

Value of the CLoPy platform in systems neuroscience

The CLoPy platform addresses a critical gap in systems neuroscience by providing a flexible and cost-effective closed-loop feedback system tailored to modulate neural activity and behavior in real time. Many current neuroscience studies assess brain activity and behavior separately, with analyses performed post hoc. CLoPy advances this paradigm by enabling researchers to directly investigate the dynamic interplay between these domains through real-time feedback, a fundamental principle of neuroplasticity and motor learning (Fetz 1969 [↗](#); Wolpert, Diedrichsen,

and Flanagan 2011 [↗](#)). This integration has broad implications for experimental neuroscience, allowing researchers to probe causal relationships between brain activity, behavior, and environmental contingencies in a controlled and scalable manner.

By enabling closed-loop experiments based on neural imaging (e.g., GCaMP fluorescence) and real-time behavioral tracking (e.g., DeepLabCut-Live), CLoPy allows researchers to address a range of experimental questions. For example, the system can be used to study how cortical plasticity facilitates motor learning, how the brain adapts to rule changes in a task, and how feedback mechanisms influence recovery from neurological injuries. These capabilities align with the needs of systems neuroscience researchers who seek to elucidate the mechanisms underlying adaptive brain-behavior interactions across different experimental conditions. Compared to other platforms, such as pyControl (Akam et al. 2022 [↗](#)), CLoPy uniquely integrates mesoscopic calcium imaging, real-time tracking, and feedback delivery, making it particularly suitable for exploring neural dynamics at a finer resolution.

Positioning within the open-source neuroscience community

As part of the broader open-source neuroscience movement, CLoPy exemplifies the principles of accessibility, modularity, and reproducibility. The platform's reliance on widely available hardware, such as Raspberry Pi and Nvidia Jetson, and its use of Python-based libraries ensure that researchers with varying levels of technical expertise can adopt and modify it. This is particularly important in the current landscape, where open-source tools like pyControl (Akam et al. 2022 [↗](#)), OpenEphys (Siegle et al. 2017 [↗](#)), and Bonsai (Lopes and Monteiro 2021 [↗](#)) have transformed experimental neuroscience by reducing costs and fostering collaboration.

CLoPy complements these efforts by extending open-source frameworks to the realm of closed-loop neuroscience. It provides detailed documentation, hardware schematics, and adaptable software, empowering researchers to customize the system for diverse experimental needs. For example, CLoPy enables experiments involving optogenetic stimulation, multimodal feedback, and integration with advanced imaging techniques. By lowering barriers to entry, it democratizes access to closed-loop experimental paradigms, particularly for labs with limited resources.

Experimental applications and implications

CLoPy's flexibility and modularity make it suitable for a wide range of experimental applications. In this study, we demonstrated its utility in closed-loop neural feedback (CLNF) and closed-loop movement feedback (CLMF) paradigms. These experiments showcase how CLoPy can be used to study fundamental questions about the dynamic interplay between neural activity, feedback, and behavior. For instance, the system's ability to deliver graded auditory feedback in response to cortical activation allowed us to explore how real-time neurofeedback modulates neural dynamics. Similarly, its real-time tracking capabilities enabled precise feedback delivery based on specific behavioral events, such as limb movements, facilitating the investigation of motor learning and adaptation. In our system, we measured latencies of ≈ 63 ms for CLNF and ≈ 67 ms for CLMF. While such latencies may limit applications requiring millisecond precision, such as fast whisker movements, saccades, or fine-reaching kinematics, we emphasize that many relevant behaviors, including postural adjustments, limb movements, locomotion, and sustained cortical state changes, occur on timescales that are well within the capture range of our system. It is also important to note that these latencies are not solely dictated by hardware constraints. A significant component arises from the inherent biological dynamics of the calcium indicator (GCaMP6s) and calcium signaling itself, which introduce slower temporal kinetics independent of processing delays. Newer variants, such as GCaMP8f, offer faster response times and could further reduce effective biological latency in future implementations. We acknowledge that Raspberry Pi provides a low-cost solution but contributes to modest computational delays, while Nvidia Jetson offers faster inference at higher cost. Our choice reflects a balance between accessibility, cost-effectiveness, and performance, making the system deployable in many laboratories. Importantly, the modular and open-source design means the pipeline can readily be adapted to higher-performance GPUs or integrated with electrophysiological recordings, which provide higher

temporal resolution. In the case of CLNF we have focused on large regional cortical GCaMP signals that are relatively slow in kinetics. While such changes are well suited for transcranial mesoscale imaging assessment, it is possible that cellular 2-photon imaging (Yu et al. 2021 [DOI](#)) or preparations that employ cleared crystal skulls (Kim et al. 2016 [DOI](#)) could resolve more localized and higher frequency kinetic signatures. In this case it would be important to potentially tune and enhance closed-loop methodology to more closely follow these events.

The platform's ability to modulate neural and behavioral outputs in real time has far-reaching implications for experimental neuroscience. It enables researchers to study neural mechanisms underlying learning, adaptation, and plasticity with a level of precision and scalability that was previously difficult to achieve. For example, CLoPy can be used to investigate how the brain reorganizes itself after injury or how neural circuits adapt to changing task demands.

Future directions and broad impact

To further enhance its utility, future iterations of CLoPy could integrate advanced machine learning algorithms for real-time data analysis and decision-making, enabling more sophisticated feedback paradigms. Additionally, expanding the platform's compatibility with other experimental modalities, such as multi-photon imaging and electrophysiology, could broaden its applications in neuroscience research. CLoPy's scalability also opens avenues for its use in translational research, such as testing neurorehabilitation interventions or brain-computer interfaces in larger animal models or humans. By providing an open-source, flexible, and cost-effective solution, CLoPy has the potential to significantly advance the field of systems neuroscience. It not only enables researchers to address complex experimental questions but also fosters collaboration and innovation within the neuroscience community. As more labs adopt and refine this platform, we anticipate that it will accelerate the pace of discovery, enhance reproducibility, and contribute to the development of novel therapeutic strategies for neurological disorders. While we aimed to include multiple animals for each experimental condition, certain task rules are represented by only one or a few animals. This limitation primarily arises from logistical challenges, including the extensive time required for training under specific task contingencies and the variability in maintaining high-quality optical windows over prolonged experimental periods. Consequently, while these data provide valuable insights into the feasibility and mechanisms of the task, we acknowledge that broader generalizations should be made cautiously and that further replication will be important in future studies.

In conclusion, our study highlights the significant potential of closed-loop feedback systems for advancing neuroscience research. By providing a flexible, cost-effective, and open-source platform, we offer a valuable tool for exploring the complex interactions between brain activity and behavior, with implications for both basic research and clinical applications.

Materials and methods

Animals

Mouse protocols were approved by the University of British Columbia Animal Care Committee (ACC) and followed the Canadian Council on Animal Care and Use guidelines (protocol A22-0054). A total of 56 mice (postnatal 104-140) were used in this study: 26 female and 30 male transgenic C57BL/6 mice expressing GCaMP6s were used. CLNF experiments (n=40, 17 females, 23 males) were done with tetO-GCaMP6s x CAMK tTA (Wekselblatt et al., 2016 [DOI](#)), and CLMF experiments (n=16, 9 females, 7 males) were done with Ai94 from the Allen Institute for Brain Science, crossed to Emx1-cre and CaMK2-tTA line (Jackson Labs) (Madisen et al., 2015). Mice were housed in a conventional facility in plastic cages with micro-isolator tops and kept on a normal 12 hr. light cycle with lights on at 7 AM. Most experiments were performed toward the end of the mouse light cycle. Mice that were unable to achieve a success rate of 70% after 7 days of training in CLNF experiments were excluded from the study (total 6 mice).

Animal surgery, chronic transcranial window preparation

Animals were anesthetized with isoflurane, and a transcranial window was installed as previously described (Silasi et al. 2016 [DOI](#); Vanni and Murphy 2014 [DOI](#)) and in an amended and more extensive protocol described here. A sterile field was created by placing a surgical drape over the previously cleaned surgical table, and surgical instruments were sterilized with a hot bead sterilizer for 20 s (Fine Science Tools; Model 18000–45). Mice were anesthetized with isoflurane (2% induction, 1.5% maintenance in air) and then mounted in a stereotactic frame with the skull level between lambda and bregma. The eyes were treated with eye lubricant (Lacrilube; www.well.ca [DOI](#)) to keep the cornea moist, and body temperature was maintained at 37°C using a feedback-regulated heating pad monitored by a rectal probe. Lidocaine (0.1 ml, 0.2%) was injected under the scalp, and mice also received a 0.5 ml subcutaneous injection of a saline solution containing buprenorphine (2 mg/ml), atropine (3 µg/ml), and glucose (20 mM). The fur on the head of the mouse (from the cerebellar plate to near the eyes) was removed using a fine battery-powered beard trimmer, and the skin was prepared with a triple scrub of 0.1% Betadine in water followed by 70% ethanol. Respiration rate and response to toe pinch were checked every 10–15 min to maintain the surgical anesthetic plane.

Before starting the surgery, a cover glass was cut with a diamond pen (Thorlabs, Newton, NJ, USA; Cat#: S90W) to the size of the final cranial window (Ø9 mm diameter). A skin flap extending over both hemispheres approximately 3 mm anterior to bregma and to the posterior end of the skull and down lateral was cut and removed. A #10 scalpel (curved) and sterile cotton tips were used to gently wipe off any fascia or connective tissue on the skull surface, making sure it was completely clear of debris and dry before proceeding. The clear version of C and B-Metabond (Parkell, Edgewood, NY, USA; Product: C and B Metabond) dental cement was prepared by mixing 1 scoop of C and B Metabond powder (Product: S399), 7 drops of C and B Metabond Quick Base (Product: S398), and one drop of C and B Universal catalyst (Product: S371) in a ceramic or glass dish (do not use plastic). Once the mixture reaches a consistency that makes it stick to the end of a wooden stir stick, a titanium fixation bar (22.2 × 2.7 × 3.2 mm) was placed so that there was a 4 mm posterior space between the bar edge and bregma, by applying a small amount of dental cement and holding it pressed against the skull until the cement partially dried (1–2 min). With the bar in place, a layer of dental adhesive was applied directly on the intact skull. The precut cover glass was gently placed on top of the mixture before it solidified (within 1 min), taking care to avoid bubble formation. If necessary, extra dental cement was applied around the edge of the cover slip to ensure that all the exposed bone was covered, and that the incision site was sealed at the edges. The skin naturally tightens itself around the craniotomy, and sutures are not necessary. The mixture remains transparent after it solidifies, and one should be able to clearly see large surface veins and arteries at the end of the procedure. Once the dental cement around the coverslip is completely solidified (up to 20 min), the animal received a second subcutaneous injection of saline (0.5 ml) with 20 mM of glucose and was allowed to recover in the home cage with an overhead heat lamp and intermittent monitoring (hourly for the first 4 hr. and every 4–8 hr. thereafter for activity level). Then, the mouse was allowed to recover for 7 days before task training.

Water deprivation and habituation to experiments

Around 10–21 days after the surgery, the animals were placed on a schedule of water deprivation. Given the variation in weight due to initial ad libitum water consumption, the mouse weight was defined 24 hr. after the start of water restriction. If mice did not progress well through training, they were still given up to 1 ml of water daily, there was also a 15% maximal weight loss criterion used for supplementation (see detailed protocol). All mice were habituated for 5 days before data collection. Awake mice were head-fixed and placed in a dark imaging chamber for training and data collection for each session. We tried to not keep animals headfixed for more than 45 minutes in each session as they become less engaged with long duration headfixed sessions. After headfixing them, it takes about 15 minutes to get the experiment going and therefore 30 - 40

minutes long recorded sessions seemed appropriate before they stop being engaged or before they get satiated in the task. High performing animals were able to maintain body weight and gain weight towards pre-surgery and pre-water restriction values.

CLNF, CLMF setup and behavior experiments

Our goal has been to deliver a robust, cross-platform, and cost-effective solution for closed-loop feedback experiments. We have designed and tested a behavioral paradigm where head-fixed mice learn an association between their cortical or behavioral activity, external feedback, and rewards. We tested our CLNF system on Raspberry Pi, and CLMF system on an Nvidia Jetson GPU device, for their compactness, general-purpose input/output (GPIO) programmability, and wide community support. While our investigations center around mesoscale cortical imaging of genetically encoded calcium sensors and behavior imaging, the approach could be adapted to any video or microscopy-dependent signal where relative changes in brightness or keypoint behavior are observed and feedback is given based on a specific rule. This system benefits from advances in pose estimation (Forys et al. 2020 [↗](#); Kane et al. 2020 [↗](#)) and employs strategies to improve real-time processing of pre-defined keypoints on compact computers such as the Nvidia Jetson. We have constructed a software and hardware-based platform built around open-source components.

To examine the performance of the closed-loop system, we used water-restricted adult transgenic mice that expressed GCaMP6s widely in the cortex (details in the methods section and in [Table 5](#) [↗](#) and [Table 6](#) [↗](#)). Using the transcranial window imaging technique, we assessed the ability of these animals to control brain regions of interest and obtain water rewards. After an initial period of habituation to manual head fixation, mice were switched to closed-loop task training. In training, we had two primary groups: 1) a single cortical ROI linked to water reward and the subject of auditory feedback ([Animation 1A, 1B, 1C](#) [↗](#)), or 2) a pairing between two cortical ROIs where auditory feedback and water rewards were given based on the difference in activity between sites ([Animation 1D, 1E, 1F](#) [↗](#)).

For neurofeedback-CLNF, we calculate an online $\Delta F/F_0$ value (GCaMP signal) which represents a relative change of intensity in the region(s) of interest. In brief, a running baseline (F_0) is computed at each sample point, subtracted and divided ($\Delta F/F_0$) from the raw fluorescence values; more details about online $\Delta F/F_0$ calculation are in the methods section. These calculations are made in near real-time on the Raspberry Pi and are used to control the GPIO output pins that provide digital signals for closed-loop feedback and water rewards to water-restricted mice. In 1-ROI experiments, we mapped the range of average $\Delta F/F_0$ values in the ROI ([Figure 1](#) [↗](#)) to a range of audio frequencies (1 kHz - 22 kHz), which acted as feedback to the animal. For 2-ROI experiments, the magnitude of the $\Delta F/F_0$ activity difference between the ROIs, based on the specified rule (e.g., ROI1-ROI2), was mapped to the range of audio frequencies. We confirmed that these frequencies were accurately generated and mapped by audio recordings ([Supplementary Figure 2](#) [↗](#) and see Methods) obtained at 200 kHz using an ultrasonic microphone (Dodotronic, Ultramic UM200K) positioned within the recording chamber ~10 cm from the audio speaker. To confirm the timing of feedback latency, LED lights were triggered instead of water rewards ([Supplementary Figure 1](#) [↗](#)), and the delay was calculated between the detected event (green LED ON for CLNF and paw movement for CLMF) and the red LED flash. In the case of CLNF, the camera recording brain activity was used to record both the flashing green and red LEDs. Temporal traces of green and red LED pixels were extracted from the recorded video, and the average delay between the green and red LEDs becoming bright was calculated as the delay in closed-loop feedback for CLNF experiments. Performing this analysis indicated that the Raspberry Pi system could provide reliable graded feedback using GPIO within 63 ± 15 ms for CLNF experiments.

An imaging rig was developed using two Raspberry Pi 4B+ single-board computers, designated as “brain-pi” (master device for initiating all recordings) for widefield GCaMP imaging and “behavior-pi” (slave device waiting for a trigger to start recording) for simultaneous behavior recording. Both devices were connected to the internet via Ethernet cables and communicated with each other through 3.3 V transistor-transistor logic (TTL) via general-purpose input outputs (GPIOs). To synchronize session initiation, GPIO pin #17 on the brain-pi (configured as an output)

was connected to GPIO pin #17 on the behavior-pi (configured as an input), allowing the brain-pi to send a TTL signal to the behavior-pi at the start of each session. Additionally, frames were aligned based on a LED ON event for each session.

Additional hardware components were integrated into the brain-pi setup. GPIO pin#27 (output) was connected to a solenoid (Gems Sensor, 45M6131) circuit to deliver water rewards, while GPIO pin#12 (output) was linked to a buzzer (Adafruit product #1739) positioned under the head-fixing apparatus to signal trial failures. GPIO pin#21 (output) was used to trigger a LED driver controlling both short blue (447.5 nm) and long blue (470 nm) LEDs, which were essential for cortical imaging. A speaker was connected to the brain-pi's 3.5 mm audio jack to provide auditory output (PulseAudio driver) during the experiments. Both the brain-pi and behavior-pi devices were equipped with compatible external hard drives, connected via USB 3.0 ports, to store imaging and behavioral data, ensuring reliable data capture throughout the experimental sessions.

Mice with implanted transcranial windows on the dorsal cortex were head-fixed in a transparent acrylic tube (1.5-inch outer diameter, 1 1/8 inch inner diameter) and placed such that the imaging camera (RGB Raspberry Pi Camera, OmniVision OV5647 CMOS sensor) was above the transcranial window, optimally focused on the cortical surface for GCaMP imaging. The GCaMP imaging cameras had lenses with a focal length of 3.6 mm and a field of view of $\approx 10.2 \times 10.2$ mm, leading to a pixel size of ≈ 40 μ m, and were equipped with triple-bandpass filters (*Chroma 69013m* [Chroma 69013m](#)), which allowed for the separation of GCaMP epifluorescence signals and 447 nm reflectance signals into the green and blue channels, respectively. The depth of field (≈ 3 mm) was similar to previous reports ([Lim et al. 2012](#) [Lim et al. 2012](#)) which provided both a large focal volume over which to collect fluorescence and made the system less sensitive to potential changes in the z-axis position. To reduce image file size, we binned data at 256×256 pixels on the camera for brain imaging data. Brain images were saved as 8-bit RGB stacks in HDF5 file format. We manually fixed the camera frame rate to 15 Hz, turned off automatic exposure and auto white balance, and set white balance gains to unity. For both green epifluorescence and blue reflection channels, we adjusted the intensity of illumination so that all values were below 180 out of 256 grey levels (higher levels increase the chance of cross-talk and saturation).

Mesoscale GCaMP imaging can often be performed with single-wavelength illumination ([Gilad and Helmchen 2020](#) [Gilad and Helmchen 2020](#); [Nakai et al. 2023](#) [Nakai et al. 2023](#)). However, in this experiment, we utilized dual-LED illumination of the cortex ([Michelson et al. 2023](#) [Michelson et al. 2023](#)). One LED (short-wavelength blue, 447 nm Royal Blue Luxeon Rebel LED SP-01-V4 paired with a Thorlabs FB 440-10 nm bandpass filter) monitored light reflectance to account for hemodynamic changes ([Xiao et al. 2017](#) [Xiao et al. 2017](#)), while the second LED (long-wavelength blue, 470 nm Luxeon Rebel LED SP-01-B6 combined with a Chroma 480/30 nm filter) was used to excite GCaMP for green epifluorescence. Both signals were captured simultaneously using an RGB camera. For each mouse, light from both the excitation and reflectance LEDs was channeled into a single liquid light guide, positioned to illuminate the cortex ([Figure 1v](#) [Figure 1v](#)) ([Michelson et al. 2023](#) [Michelson et al. 2023](#)). Further specifics are outlined in the accompanying Parts List and assembly instructions document. A custom-built LED driver, controlled by a Raspberry Pi, activated each LED at the beginning of the session and deactivated them at the session's end. This on-off illumination shift was later used in post hoc analysis to synchronize frames from the brain and behavior cameras.

Before initiating the experimental session, key session parameters were configured in the config.ini file, as detailed in the "Key Configuration Parameters" section below. Additionally, we ensured that the waterspout was positioned appropriately for the mouse to access and consume the reward. To launch the experiment, open a terminal on the brain-pi (master) device, ensuring that the current directory is set to clcopy/brain/. The experiment can be started by entering the command `python3 <SCRIPT_NAME>.py`, where <SCRIPT_NAME> corresponds to the appropriate Python script for the session ([Supplementary Figure 7](#) [Supplementary Figure 7](#)). We provide two pre-defined scripts in the codebase: one for 1ROI and another for 2ROI experiments. Both scripts are functionally similar. Upon execution, the script prompts for the "mouse_id," which can be entered as an array of characters, followed by pressing 'Enter.' Upon initialization, two preview windows appear. The first window displays live captured images with intensity values overlaid in the green and blue

channels, which allow for adjustments to LED brightness levels. The second window displays real-time $\Delta F/F_0$ values overlaid on cortical locations, enabling brain alignment checks. These preview windows are used to assess imaging quality and ensure appropriate settings before starting the experiment (Supplementary Figure 6 [↗](#)). Once all settings are confirmed, pressing the ‘Esc’ key starts the experiment. At this point, only one window displaying incoming images is shown. The experiment begins with a rest period of duration specified in the config.ini file, followed by alternating trial and rest periods. Data acquired during the session are saved in real-time on the device. For more details and the latest instructions on system setup, please refer to the project’s GitHub page.

An additional imaging rig for CLMF was developed utilizing an Nvidia-Jetson Orin device (8-core ARM Cortex CPU, 2048 CUDA cores, 64 GB memory), which served as the master device responsible for triggering all recordings. The Nvidia-Jetson was connected to an Omron Sentech STC-MCCM401U3V USB3 Vision camera for behavior imaging and real-time pose tracking. A personal computer (serving as the slave device) running EPIX software and frame grabber (PIXCI® E4 Camera Link Frame Grabber) was connected to a Pantera TF 1M60 CCD camera (Dalsa) for widefield GCaMP imaging. Both the Nvidia-Jetson and the EPIX PC were linked via transistor-transistor logic (TTL) for communication. For session synchronization, the GPIO pin #17 on the Nvidia-Jetson, configured as an output, was connected to the trigger input pin of the EPIX framegrabber E4DB. This configuration enabled the Nvidia-Jetson to send a TTL signal to initiate frame capture on the EPIX system at the start of each session. Additionally, this same pin was connected to an LED driver to trigger the activation of a long blue LED (470 nm) for GCaMP imaging simultaneously.

Several hardware components were integrated into the Nvidia-Jetson setup to support experimental protocols. GPIO pin #13 (output) was linked to a solenoid circuit to deliver water rewards, while GPIO pin #7 (output) was connected to a buzzer placed under the head-fixation apparatus to signal trial failures. Auditory feedback during the experiments was provided through a speaker connected to the Nvidia-Jetson’s audio output. Both the Nvidia-Jetson and EPIX PC were equipped with sufficient storage space to ensure the reliable capture and storage of behavioral video recordings and widefield image stacks, respectively, throughout the experimental sessions.

Mice, implanted with a transcranial window over the dorsal cortex, were head-fixed in a custom-made transparent, rectangular chamber (dimensions provided). The chamber was designed using CAD software and fabricated from a 3 mm-thick acrylic sheet via laser cutting. Complete model files, acrylic sheet specifications, and assembly instructions are available on the GitHub repository. A mirror was positioned at the bottom and front of the chamber (Figure 1B [↗](#)) to allow multiple views of the mouse for improved tracking accuracy. The Omron Sentech camera was equipped with an infrared (IR)-only filter to exclusively capture IR illumination, effectively blocking other light sources and ensuring consistent behavioral imaging for accurate pose tracking. For widefield GCaMP imaging, the Dalsa camera was equipped with two front-to-front lenses (50 mm, f/1.435 mm and f/2; Nikon Nikkor) and a bandpass emission filter (525/36 nm, Chroma). The 12-bit images were captured at a frame rate of 30 Hz (exposure time of 33.3 ms) with 8×8 on-chip spatial binning (resulting in 128×128 pixels) using EPIX XCAP v3.8 imaging software. The cortex was illuminated using a blue LED (473 nm, Thorlabs) with a bandpass filter (467–499 nm) to excite calcium indicators, and the blue LED was synchronized with frame acquisition via TTL signaling. GCaMP fluorescence image stacks were automatically saved to disk upon session completion.

Prior to starting a session, key experimental parameters were configured within the config.ini file (see “Key Configuration Parameters” below for details). The waterspout was positioned close to the mouse to allow easy access for licking and consuming rewards. To initiate the experiment on the Nvidia-Jetson (master device), the terminal’s current directory was set to `clcopy/behavior/`. The experiment was launched by executing the command `python3 <SCRIPT_NAME>.py`, where `<SCRIPT_NAME>` corresponds to the Python script responsible for running the experiment (Supplementary Figure 6 [↗](#)). The provided script maps the speed of a specific control-point (FLL_bottom) to the audio feedback and executes 60 trials. Once launched, the program prompts for the “mouse_id,” which could be entered as an array of characters, followed by pressing Enter.

A preview window displaying a live behavioral view was then presented, allowing for adjustments of IR brightness levels and the waterspout (Supplementary Figure 6 [↗](#)). Once imaging quality and settings were confirmed to be optimal, pressing the “esc” key started the experiment. During the experiment, the tracked points were overlaid on the real-time video, and the session alternated between rest and trial periods. All acquired data were saved locally on the Nvidia-Jetson during the experiment and automatically transferred to the EPIX PC after completion. Additional setup instructions and system details can be found on the associated GitHub page.

Key configuration parameters

Prior to running the experiment script, session parameters were set in the config.ini file under a designated configuration section. This file contains multiple config sections, each tailored to a different experiment type, such as brain-pi or behavior-pi. The appropriate config section is specified in the experiment script. Key parameters in this section are listed in Table 1 [↗](#).

Further details on these parameters can be accessed via the GitHub repository.

CLoPy platform

Closed-Loop Feedback Training System (CLoPy) is an open-source software and hardware system implemented in the Python (≥ 3.8) programming language. This work is accompanied by a package to replicate the system, reproduce figures in this publication, and an extensive supplemental guide with full construction illustrations and parts lists to build the platform used. See <https://github.com/pankajkgupta/clopy> [↗](#) for details and accompanying acquisition and analysis code. We also provide model files for machined and 3D-printed parts in the repository; links to neural and behavioral data can also be found at the federated research data repository-<https://doi.org/10.20383/102.0400> [↗](#).

While the presented CLNF experiments were conducted on a Raspberry Pi 4B+ device, the system can be used on any other platform where the Python runtime environment is supported. Similarly, CLMF experiments were conducted on an Nvidia Jetson Orin device, but it can be deployed on any other device with a GPU for real-time inference. For any programmable camera to be used with the system, one can implement a wrapper Python class that implements the CameraFactory interface functions for integration with the system.

Graded feedback (online $\Delta F/F_0$ -audio mapping)

The auditory feedback was proportional to the magnitude of the neural activity. We translated fluorescence levels into the appropriate feedback frequency and played the frequency on speakers mounted on two sides of the imaging platform. Frequencies used for auditory feedback ranged from 1 to 22 kHz in quarter-octave increments (Han et al. 2007 [↗](#); Clancy et al. 2014 [↗](#)). When a target was hit, a circuit driven solenoid delivered a reward to mice.

To map the fluorescence activity F to the quarter-octave index n , we used linear scaling. Let F_{min} and F_{max} be the minimum and maximum fluorescence values n_{max} . The maximum number of quarter-octave steps between 1 kHz and 22 kHz.

The number of quarter-octave steps between 1 kHz and 22 kHz can be calculated as:

$$n_{max} = 4 \cdot \log_2 \left(\frac{f_{max}}{f_{min}} \right)$$

where $f_{max} = 22\text{kHz}$ and $f_{min} = 1\text{kHz}$

The quarter-octave index n can be mapped linearly from fluorescence activity F as:

$$n = \frac{F - F_{min}}{F_{max} - F_{min}} \cdot n_{max}$$

Thus, the final equation mapping fluorescence activity to audio frequency is:

Parameter	Description
vid_source	Specifies the class responsible for the image stream, which may come from a programmable camera or another video source.
data_root	Directory path for saving the recorded sessions and current

Table 1. Key configuration parameters in CLoPy.

	configuration.
raw_image_file	File name for saving the image stream.
resolution	Image stream resolution in (x, y) pixels.
framerate	Number of frames per second from the image stream.
awb_mode	Auto-white-balance mode (only used in CLNF, True for Behavior-Pi, False for Brain-Pi).
shutter_speed	Sets the camera sensor exposure.
dff_history	The duration (in seconds) used to calculate $\Delta F/F_0$, only used in CLNF.
ppmm	Pixels-per-mm value at the focal plane of the camera.
bregma	Y, X pixel coordinates of the bregma on the dorsal cortex in the image frame, only used in CLNF.
seeds_mm	List of cortical locations for inclusion in the closed-loop training rule, each defined by a name and coordinates relative to bregma (in mm), only used in CLNF.
roi_operation	Specifies the ROI(s) and operation (+ or -) for closed-loop training.
roi_size	Size of the ROI(s) in mm.
n_tones	Number of distinct audio frequencies for graded auditory feedback.
reward_delay	Delay (in seconds) after crossing the threshold.
reward_threshold	Threshold value in terms of $\Delta F/F_0$, determined from baseline sessions.
adaptive_threshold	A setting for adjusting the reward threshold. If set to 0, the threshold remains constant throughout the session; if set to 1, the threshold increases or decreases by 0.02 steps based on the reward rate.
total_trials	Total number of trials in the session.
max_trial_dur	Maximum trial duration (in seconds).
success_rest_dur	Rest duration after a successful trial (in seconds).
fail_rest_dur	Rest duration after a failed trial (in seconds).
initial_rest_dur	Rest period at the beginning of the session before the first trial (in seconds).
summary_file	File name to save the experiment summary as comma-separated values (CSV).

Table 1. (continued)

summary_header	List of variable names to be saved in the summary file.
dlc_model_path	Path of the DeepLabCut-Live model for real time pose tracking (used only in CLMF).
control_point	Name of the tracked point for closed-loop feedback (used only in CLMF).
speed_threshold	Speed threshold for a success in a trial and receive reward (used only in CLMF).

Table 1. (continued)

$$f(F) = f_{min} \cdot 2^{\left(\frac{1}{4} \cdot \frac{F - F_{min}}{F_{max} - F_{min}} \cdot n_{max} \right)}$$

where $f_{min} = 1\text{kHz}$, $f_{max} = 22\text{kHz}$, and $n_{max} = 4 \cdot \log_2(22) \approx 13.45$

This equation allows us to map any fluorescence activity F within the range $[F_{min}, F_{max}]$ to a corresponding frequency within the 1–22 kHz range, in quarter-octave increments.

Checking the dynamic range of graded auditory feedback

To assess the dynamic range of audio signals from the speakers, we used an Ultra microphone (Dodotronic) to record feedback signals generated during a session, detect the frequencies, and compare the detected frequencies through our speakers to what was commanded through the program. We verified a linear relationship between 1 and 22 kilohertz (Supplementary figure 2 [↗](#)). Analysis of GCAMP imaging experiments indicated that baseline activity values were associated with mapped sounds in the 6±4 kHz range. At reward points, which are threshold-dependent ($\Delta F/F$ over F values), the commanded auditory feedback values were significantly higher (18±2 kHz range). A previous version of the CLNF system was found to have non-linear audio generation above 10 kHz, partly due to problems in the audio generation library and partly due to the consumer-grade speaker hardware we were employing. This was fixed by switching to the Audiostream (<https://github.com/kivy/audiostream> [↗](#)) library for audio generation and testing the speakers to make sure they could output the commanded frequencies (supplementary figure 2 [↗](#)).

Determining reward threshold based on a baseline session

Before starting the experiments, and after the habituation period, a baseline session (day 0) of the same duration as the experiments is recorded. This session is similar to the experimental sessions in the coming days, except that the mice do not receive any rewards. Offline analysis of this session is used to establish a threshold value for the target activity (see the target activity section to read more). In brief, for 1 ROI experiments, target activity is the average $\Delta F/F_0$ activity in that ROI. For 2 ROI experiments, target activity is based on the specified rule in the config.ini file. For example, target activity for the rule “ROI1-ROI2” would be “average $\Delta F/F_0$ activity in ROI1 - average $\Delta F/F_0$ activity in ROI2.”

Determining ROI(s) for changes in CLNF task rules

Dorsal cortical widefield activity is dynamic, and ongoing spatiotemporal motifs involve multiple regions changing activity. Choosing ROI(s) for CLNF has some caveats and requires some considerations before choosing. It was relatively straightforward for initial training experiments where baseline (day 0) sessions were used to establish a threshold for the selected ROI on day 1. Mice learn to modulate the selected ROI over the training sessions. Interestingly, other cortical ROIs were also changing along with the target ROI as mice were learning the task. We needed to be careful when changing the ROI on day 11 because if we changed to an ROI that also changes along with the previous ROI, mice could keep getting rewards without realizing any change. To address this issue, we analyzed the neural data from day 1 to day 10 and found potential ROIs for which the threshold crossings did not increase significantly or were not on par with the previous ROI, and one of these ROIs was chosen for the rule change.

Online $\Delta F/F_0$ calculation

In calcium imaging, $\Delta F/F_0$ is often used to represent the change in fluorescence relative to a baseline fluorescence (F_0), which helps in normalizing the data. For CLNF real-time feedback, we computed $\Delta F/F_0$ online, frame by frame.

Let $F(t)$ represent the fluorescence signal at time t . A running baseline $F_0(t)$ is estimated by applying a sliding window to the fluorescence signal to calculate a moving average over a window of size N .

$$F_0(t) = \frac{1}{N} \sum_{i=t-N}^t F(i)$$

Once the running baseline $F_0(t)$ is computed, the $\Delta F/F_0$ at time t is calculated as:

$$\frac{\Delta F(t)}{F_0(t)} = \frac{F(t) - F_0(t)}{F_0(t)}$$

Each incoming frame contained both a green channel (capturing GCaMP6s fluorescence) and a short blue channel (representing blood volume reflectance) (Ma et al. 2016 [↗](#); Wekselblatt et al. 2016 [↗](#); Valley et al. 2020 [↗](#)). These frames were appended to a doubly-ended queue (deque), a Python data structure, for running baseline F_0 (□) The length of the deque (□) was defined by the product of two parameters from the configuration file: “dff_history” (in seconds) and “framerate” (in frames per second), which determined the number of frames used for the running baseline. For the CLNF experiments, we used a dff_history of 6 s (a dynamic baseline) and a framerate of 15 fps, resulting in a deque length of 90 frames. When the deque reached full capacity, appending a new frame automatically removed the oldest frame, ensuring an updated running $\Delta F/F_0$ throughout the experiment, regardless of session duration. This method effectively avoided memory limitations over time. We would like to emphasize that we are not directly averaging activity over 6 s to compare against the reward threshold. Instead, the preceding 6 s of activity is used solely to compute a dynamic baseline for $\Delta F/F_0$ ($\Delta F/F_0 = (F - F_0)/F_0$). Here, F_0 is calculated as the mean fluorescence intensity over the prior 6 s window and is updated continuously throughout the session. This baseline is then subtracted from the instantaneous fluorescence signal to detect relative changes in activity. The reward threshold is therefore evaluated against these baseline-corrected $\Delta F/F_0$ values at the current time point, not against an average over 6 s. This moving-window baseline correction is a standard approach in calcium imaging analyses, as it helps control for slow drifts in signal intensity, bleaching effects, or ongoing fluctuations unrelated to the behavior of interest. Thus, the 6 s window is not introducing a temporal lag in reward assignment but is instead providing a reference to detect rapid increases in cortical activity. For each channel (green and blue), the running F_0 was subtracted from each incoming frame to obtain ΔF , followed by $\Delta F/F_0$ calculations, applied per pixel as a vectorized operation in Python. To correct for hemodynamic artifacts (Ma et al. 2016 [↗](#)), we subtracted the blue excitation and epifluorescence channel $\Delta F/F_0$ from the green channel $\Delta F/F_0$.

It is important to note that while blood volume reflectance is typically captured using green light (Ma et al. 2016 [↗](#)), we used short blue light due to technical constraints associated with the Raspberry Pi camera’s rolling shutter which made strobing infeasible. The short blue light (447 nm) with a 440 ± 5 nm filter is close to the hemoglobin isosbestic point and has been shown to correlate well with the 530 nm green light signal as a proxy for hemodynamic activity (Michelson et al. 2023 [↗](#); Murphy et al. 2020 [↗](#)). Additionally, the 447 nm LED would be expected to produce minimal green epifluorescence at the low power settings used in our experiments (Dana et al. 2014 [↗](#)). Previous studies have evaluated and compared the performance of corrected versus uncorrected signals using this method (Xiao et al. 2017 [↗](#); Murphy et al. 2020 [↗](#)).

Offline $\Delta F/F_0$ calculation

For offline analysis of GCaMP6s fluorescence in CLNF experiments, the green and blue channels (Ma et al. 2016 [↗](#); Wekselblatt et al. 2016 [↗](#); Valley et al. 2020 [↗](#)) were converted to $\Delta F/F_0$ values. For each channel, a baseline image (F_0) was computed by averaging across all frames of the recording session. The F_0 was then subtracted from each individual frame producing a difference image (ΔF). This difference was divided by F_0 , resulting in the fractional change in intensity ($\Delta F/F_0$) for each pixel as a function of time. To further correct for hemodynamic artifacts (Ma et al., 2016 [↗](#)), the blue channel reflected light $\Delta F/F_0$ signal, reflecting blood volume changes, was subtracted from the green channel $\Delta F/F_0$ signal, isolating the corrected GCaMP6s fluorescence response from any potential confounding vascular contributions.

Offline analysis of GCaMP6s fluorescence in CLMF experiments involved processing the green epifluorescence channel. We did not collect the hemodynamic signal in this experiment because we intended to employ the second channel for optogenetic stimulation (not included in this article). A baseline image (F_0) was computed by averaging across all frames of the recording session. The F_0 was then subtracted from each individual frame, producing a difference image (ΔF). This difference was divided by F_0 , resulting in the fractional change in intensity ($\Delta F/F_0$) for each pixel as a function of time.

Seed-pixel correlation matrices

Widefield cortical image stacks were registered to the Allen Mouse Brain Atlas (Wang et al. 2020 [↗](#); Chon et al. 2019 [↗](#)) and segmented into distinct regions, including the olfactory bulb (OB), anterolateral motor cortex (ALM), primary motor cortex (M1), secondary motor cortex (M2), sensory forelimb (FL), sensory hind limb (HL), barrel cortex (BC), primary visual cortex (V1), and retrosplenial cortex (RS) in both the left and right cortical hemispheres. The average activity in a $0.4 \times 0.4 \text{ mm}^2$ area (equivalent to a 10×10 pixel region) centered on these regions (also referred to as seed pixels) was calculated, representing the activity within each region. These signals were then used to generate temporal plots. The temporal plots were epoched into trial and rest conditions, and correlations between the regions were computed, resulting in correlation matrices for each condition across each day. Each element of the matrix represents the pairwise correlation between two cortical regions. For each pair of cortical regions, we obtained correlation values during both trial and rest periods for every session, creating a time series (over sessions) with two conditions (trial and rest).

Statistics

Various statistical tests were performed to support the analysis presented in accompanying figures. For p-value matrices in Figure 10B [↗](#), 10C [↗](#), a two-way repeated-measures ANOVA with Bonferroni post-hoc correction was used to test the significance of correlation changes across two factors: sessions (days 1–10) and condition (trial vs. rest). Significant changes in the correlation matrices along these two variables showed complementary patterns (Figure 10 [↗](#)). Changes across sessions involved the bilateral M1, M2, FL, HL, and BC regions, while changes between the trial and rest conditions involved bilateral OB, ALM, V1, and RS regions.

A list of statistical tests used in a figure, its purpose and data used are summarized in Table 2 [↗](#).

CLNF rules

The task rules for CLNF experiments involved cortical activity at selected locations. These locations were standard Allen Mouse Brain Atlas (Wang et al. 2020 [↗](#); Chon et al. 2019 [↗](#)) CCF coordinates. Table 3 [↗](#) lists all the task rules we tested and what they meant in the context of CLNF experiments.

CLMF rules

The task rules for CLMF experiments involved points on body parts that were tracked (listed in results section). Table 4 [↗](#) lists the task rules (i.e. tracked points) we tested and what they meant in the context of CLMF experiments.

Data availability

The source data used in this paper is available here- <https://doi.org/10.20383/103.01152> [↗](#). Code to replicate the system, recreate the figures, and associated pre-processed data are publicly available and hosted on GitHub - <https://github.com/pankajkgupta/clopy> [↗](#). Any additional information required to reanalyze the data reported in this work is available from the lead contact upon request.

Figure	Test name	Purpose	Arguments
Figure 4A	Mann-Kendall Test (pymannkendall package in Python)	To determine a monotonic upward trend in the success rate of all CLNF mice before the rule change (from day 1 to day 10).	Normalized success rate per mouse before the rule change (N=40)
		To determine a monotonic upward trend in the success rate of CLNF rule-change mice after the rule change (day 11 to day 19).	Normalized success rate per mouse after the rule change (N=17)

Table 2. List of statistical tests performed.

Figure 4A	RM-ANOVA (Pingouin package in Python)	Test the significance of the success rate increase before the rule change in the CLNF No-rule-change group (Day 1 to Day 10).	Normalized success rate per mouse in No-rule-change mice before the rule change (N=23)
		Test the significance of the success rate increase after the rule change in the CLNF Rule-change group (day 11 to day 19).	Normalized success rate per mouse in Rule-change mice after the rule change (N=17)
Figure 4A	Two way RM-ANOVA (Pingouin package in Python)	Test significance of day and group on the CLNF success rate before rule change (day 1 to day 10)	Normalized success rate per mouse in Rule-change and No-rule-change mice before rule change, along with group labels.
Figure 4A	Mann-Whitney U-test with Benjamini-Hochberg correction (pymannkendall package in Python)	Test the difference in the distribution of CLNF success rates between: day 1 no-rule-change and day 3 no-rule-change day 10 No-rule change and day 10 Rule change day 10 rule change and day 11 rule change day 10 rule change and day 19 rule change	Normalized success rate per mouse in Rule-change and No-rule-change mice before rule change, along with group labels.
Figure 4B	RM-ANOVA (Pingouin package in Python)	Test the significance of the success rate increase before the rule change in the CLMF Rule-change group (Day 1 to Day 4). Test the significance of the success rate increase after the rule change in the	Normalized success rate per mouse in Rule-change and No-rule-change mice before rule change, along with group labels.

Table 2. (continued)

		CLMF Rule-change group (from day 5 to day 10). Test the significance of the success rate increase in the CLMF no-rule-change group (day 1 to day 10). Test significance of success rate increase in CLMF No-feedback group (day 1 to day 10).	
Figure 5C	Permutation Test (MMD function in hyppo package)	Multivariate two-sample test to compare ROI1 vs ROI2 bivariate distribution on day 9 and day 19	Bivariate distribution of ROI1 and ROI2 on day 9 and day 19
Figure 9C	Mann-Whitney, corrected for multiple comparisons with Benjamini-Hochberg (pymannkendall package in Python)	Test the significance of the difference in $\Delta F/F_{0peak}$ distribution between. Day 1 and day 4 in cortical OBL of CLMF rule-change mice. day 4 and day 5 in cortical OBL of CLMF rule-change mice Day 1 and Day 4 in cortical FLL of CLMF rule-change mice. day 4 and day 5 in cortical FLL of CLMF rule-change mice Day 1 and Day 4 in cortical VIL of CLMF rule-change mice.	$\Delta F/F_{0peak}$ values in named cortical regions on each day for all mice along with their group labels.

Table 2. (continued)

		day 4 and day 5 in cortical V1L of CLMF rule-change mice	
Figure 10B, 10C	Two-way RM-ANOVA, Bonferroni corrected for multiple comparisons (Pingouin package in Python)	Test significant changes in seed pixel correlations over days and between trial and rest conditions for CLMF No-rule-change mice (Figure 10B) and for CLMF No-feedback mice (Figure 10C).	Pairwise seed pixel correlation values for all combination of cortical seed locations over days for each mouse, along with group labels.
Supplementary Figure 3A	Mann-Whitney U-test with Benjamini-Hochberg correction (pymannkendall package in Python)	Test the difference in the distribution of CLNF success rates between. day 1 male and day 3 male. day 4 male and day 4 female. day 10 female and day 11 female	Normalized success rate per mouse along with their sex labels
Supplementary Figure 3B	Mann-Whitney U-test with Benjamini-Hochberg correction (pymannkendall package in Python)	Test the difference in the distribution of CLNF success rates between. day 1 single ROI and day 3 single ROI day 10 single ROI and day 10 dual ROI	Normalized success rate per mouse along with their ROI-type of the experiment.

Table 2. (continued)

Table 3. List of CLNF task rules with their description.

CLNF rule name	Description
M1_L	$\Delta F/F_0$ activity in left hemisphere primary motor area
M1_R	$\Delta F/F_0$ activity in right hemisphere primary motor area
M2_L	$\Delta F/F_0$ activity in left hemisphere secondary motor area
HL_L	$\Delta F/F_0$ activity in left hemisphere sensory hindlimb area
RL_L	$\Delta F/F_0$ activity in left hemisphere rostralateral area of visual cortex
RL_R	$\Delta F/F_0$ activity in right hemisphere rostralateral area of visual cortex
V1_L	$\Delta F/F_0$ activity in left hemisphere primary visual area
BC_L – BC_R	$\Delta F/F_0$ activity in left hemisphere barrel cortex subtracted by $\Delta F/F_0$ activity in right hemisphere barrel cortex
HL_L – HL_R	$\Delta F/F_0$ activity in left sensory hindlimb area subtracted by $\Delta F/F_0$ activity in right hemisphere sensory hindlimb area
BC_L – HL_L	$\Delta F/F_0$ activity in left hemisphere barrel cortex subtracted by $\Delta F/F_0$ activity in left hemisphere sensory hindlimb area
M1_L – HL_L	$\Delta F/F_0$ activity in left hemisphere primary motor area subtracted by $\Delta F/F_0$ activity in left hemisphere sensory hindlimb area
TR_L - M2_L	$\Delta F/F_0$ activity in the left hemisphere rostral part of the temporal association area is subtracted by $\Delta F/F_0$ activity in the left hemisphere secondary motor area.
V1_R - V1_L	$\Delta F/F_0$ activity in the right hemisphere primary visual area subtracted by $\Delta F/F_0$ activity in the left hemisphere primary visual area.
M2_L - TR_L	$\Delta F/F_0$ activity in left hemisphere secondary motor area subtracted by $\Delta F/F_0$ activity in left hemisphere rostral part of temporal association area
HL_L - M2_L	$\Delta F/F_0$ activity in left hemisphere sensory hindlimb area subtracted by $\Delta F/F_0$ activity in left hemisphere secondary motor area

Table 4. List of CLMF task rules with their description.

CLMF rule name	Description
FLL	Speed of left forelimb in the bottom view of the camera
FLR	Speed of right forelimb in the bottom view of the camera

Table 5. List of mice – CLNF.

	MouseID	group	ROIs	initial_rule	change_rule	DoB	Expt. Start date	Age(days)	Sex	Reward rate				genotype
										Day1	Day5	Day11	Day16	
1	BM2	G1	2ROI	BC L – BC R	-	20220122	20220713	172	m	0.01	0.8	-	-	tta-gcamp6s
2	BM3	G1	2ROI	BC L – BC R	-	20220122	20220712	171	m	0.28	1	-	-	tta-gcamp6s
3	BM4	G1	2ROI	BC L – BC R	-	20220122	20220712	171	m	0.01	0.83	-	-	tta-gcamp6s
4	BG2	G1	2ROI	HL L – HL R	-	20210601	20211023	144	m	0.00	0.63	-	-	tta-gcamp6s
5	BG3	G1	2ROI	HL L – HL R	-	20210601	20211023	144	m	0.00	0.64	-	-	tta-gcamp6s
6	AC4	G1	2ROI	HL L – HL R	-	20210610	20211009	121	f	0.00	0.71	-	-	tta-gcamp6s
7	TW1	G1	1ROI	RL L	-	20200729	20210113	168	m	0.00	0.35	-	-	tta-gcamp6s
8	FM1	G1	1ROI	M1 L	-	20190915	20200129	136	f	0.13	0.66	-	-	tta-gcamp6s
9	FM2	G1	1ROI	M1 L	-	20190915	20200129	136	f	0.03	1	-	-	tta-gcamp6s
10	FM3	G1	1ROI	M1 L	-	20190915	20200129	136	f	0.02	1	-	-	tta-gcamp6s
11	FM4	G1	1ROI	M1 L	-	20190915	20200129	136	f	0.25	1	-	-	tta-gcamp6s
12	CV2	G1	2ROI	BC L – HL L	-	20190815	20191205	112	m	0.01	0.72	-	-	tta-gcamp6s
13	CV3	G1	2ROI	BC L – HL L	-	20190815	20191205	112	m	0.01	1	-	-	tta-gcamp6s
14	DU1	G1	2ROI	M1 L – HL L	-	20190719	20191205	139	f	0.01	1	-	-	tta-gcamp6s
15	DU3	G1	2ROI	M1 L – HL L	-	20190719	20191205	139	f	0.02	0.39	-	-	tta-gcamp6s
16	BZ1	G1	2ROI	TR L – M2 L	-	20190406	20190806	122	m	0.24	0.29	-	-	tta-gcamp6s
17	BZ2	G1	2ROI	TR L – M2 L	-	20190406	20190806	122	m	0.02	0.32	-	-	tta-gcamp6s
18	DU1	G1	2ROI	V1 R – V1 L	-	20190217	20190807	171	f	0.00	1	-	-	tta-gcamp6s
19	DU2	G1	2ROI	V1 R – V1 L	-	20190217	20190807	171	f	0.00	0.46	-	-	tta-gcamp6s
20	BZ1	G1	2ROI	M2 L – TR L	-	20181230	20190516	137	m	0.19	0.19	-	-	tta-gcamp6s
21	BZ2	G1	2ROI	M2 L – TR L	-	20181230	20190516	137	m	0.02	0.71	-	-	tta-gcamp6s
22	BU2	G1	2ROI	M2 L – TR L	-	20181219	20190517	149	m	0.20	1	-	-	tta-gcamp6s
23	BU3	G1	2ROI	M2 L – TR L	-	20181219	20190517	149	m	0.01	0.61	-	-	tta-gcamp6s
24	BM1	G2	2ROI	BC L – BC R	RFL L – CFLL (Day11)	20220122	20220712	171	m	0.00	0.63	0.03	0.52	tta-gcamp6s
25	AC1	G2	2ROI	BC L – BC R	HL L – HL R (Day11)	20210610	20211009	121	f	0.27	0.41	0.12	1	tta-gcamp6s
26	AC2	G2	2ROI	BC L – BC R	HL L – HL R (Day11)	20210610	20211009	121	f	0.00	1	0.29	0.96	tta-gcamp6s
27	TW1	G2	1ROI	RL L	HL L (Day11)	20200802	20210113	129	m	0.01	0.58	0.00	0.71	tta-gcamp6s
28	TW2	G2	1ROI	RL L	PM L (Day11)	20200802	20210113	129	m	0.02	0.28	0.32	0.90	tta-gcamp6s
29	TW22	G2	1ROI	RL R	RL L (Day11)	20200730	20201119	112	m	0.10	0.48	0.35	0.60	tta-gcamp6s
30	TW11	G2	1ROI	M2 L	RL L (Day11)	20200730	20201119	112	m	0.05	0.32	0.10	0.64	tta-gcamp6s
31	TW33	G2	1ROI	M1 R	RL R (Day11)	20200730	20201119	112	m	0.02	0.41	0.40	0.60	tta-gcamp6s
32	FM3	G2	1ROI	V1 L	M1 R (Day11)	20190917	20200129	134	f	0.08	0.62	0.03	0.37	tta-gcamp6s
33	FM1	G2	1ROI	M1 L	V1 L (Day11)	20190917	20200129	134	f	0.00	0.65	0.02	0.61	tta-gcamp6s
34	FM2	G2	1ROI	M1 L	M1 R (Day11)	20190917	20200129	134	f	0.10	1	0.14	1	tta-gcamp6s
35	FM3	G2	1ROI	M1 L	V1 L (Day11)	20190917	20200129	134	f	0.03	0.63	0.20	0.46	tta-gcamp6s
36	FM4	G2	1ROI	M1 L	V1 L (Day11)	20190917	20200129	134	f	0.01	0.31	0.00	0.62	tta-gcamp6s
37	DR2	G2	2ROI	M2 L – TR L	HL L – M2 L (Day11)	20190401	20190806	127	m	0.09	1	0.01	1	tta-gcamp6s
38	DR3	G2	2ROI	M2 L – TR L	HL L – M2 L (Day11)	20190401	20190806	127	m	0.03	1	0.06	0.74	tta-gcamp6s
39	DR4	G2	2ROI	M2 L – TR L	HL L – M2 L (Day11)	20190401	20190806	127	m	0.00	0.64	0.02	1	tta-gcamp6s
40	DU3	G2	2ROI	M2 L – TR L	V1 R – V1 L (Day11)	20190408	20190807	121	f	0.13	0.36	0.10	1	tta-gcamp6s

Table 6. List of mice – CLMF.

	MouseID	Group	initial_rule	change_rule	DoB	Exp. start	Age(days)	Sex	Reward rate				genotype
									Day1	Day4	Day5	Day10	
1	FW2	Rule change	FLL	FLR (Day5)	20230209	20230713	154	f	0.30	0.60	0.21	0.95	Ai94
2	FW3	Rule change	FLL	FLR (Day5)	20230209	20230718	159	f	0.09	0.65	0.00	0.92	Ai94
3	FW22	Rule change	FLL	FLR (Day5)	20230209	20230719	160	f	0.25	0.70	0.44	1	Ai94
4	GT3	Rule change	FLL	FLR (Day5)	20230219	20230715	146	f	0.11	0.93	0.63	0.85	Ai94
5	GT33	Rule change	FLL	FLR (Day5)	20230219	20230719	150	f	0.12	0.78	0.00	0.87	Ai94
6	HA2	Rule change	FLL	FLR (Day5)	20230225	20230719	144	m	0.12	0.72	0.28	0.9	Ai94
7	GER2	Rule change	FLL	FLR (Day5)	20230823	20231228	127	m	0.34	0.95	0.50	0.83	Ai94
8	HYL3	Rule change	FLL	FLR (Day5)	20230827	20240103	129	m	0.33	0.93	0.60	0.91	Ai94
9	BR1	No rule change	FLL	-	20230729	20231201	125	m	0.11	0.58	0.97	1	Ai94
10	BR2	No rule change	FLL	-	20230729	20231201	125	m	0.38	1	0.84	0.82	Ai94
11	GER1	No rule change	FLL	-	20230811	20231228	139	f	0.35	0.88	0.84	0.82	Ai94
12	HYL2	No rule change	FLL	-	20240103	20240103	121	f	0.33	0.93	0.88	0.93	Ai94
13	GIL2	No feedback	FLL	-	20231219	20231219	112	m	0.09	0.52	0.28	0.95	Ai94
14	GIL3	No feedback	FLL	-	20231219	20231219	135	f	0.23	0.53	0.57	0.51	Ai94
15	GIL4	No feedback	FLL	-	20231219	20231219	135	f	0.19	0.63	0.60	0.69	Ai94
16	GER3	No feedback	FLL	-	20230811	20231219	130	m	0.08	0.48	0.43	0.75	Ai94

Acknowledgements

This work was supported by Canadian Institutes of Health Research (CIHR) foundation grant FDN-143209 and project grant PJT-180631 (to T.H.M.). T.H.M. was also supported by the Brain Canada Neurophotonics Platform, a Heart and Stroke Foundation of Canada grant in aid, the National Science and Engineering Council of Canada (NSERC; GPIN-2022-03723), and a Leducq Foundation grant. This work was supported by resources made available through the Dynamic Brain Circuits cluster and the NeuroImaging and NeuroComputation Centre at the UBC Djavad Mowafaghian Centre for Brain Health (RRID SCR_019086) and made use of the DataBinge forum, and computational resources and services provided by Advanced Research Computing (ARC) at the University of British Columbia. We thank Pumin Wang and Cindy Jiang for surgical assistance, Jamie Boyd and Jeffrey M LeDue for technical assistance.

Additional files

[Animations](#) 

[Supplementary Figures](#) 

Additional information

Funding

Funder	Grant reference number
Canadian Institutes of Health Research (CIHR)	FDN-143209
Canadian Institutes of Health Research (CIHR)	PJT-180631

Author ORCID iDs

Pankaj K Gupta: <https://orcid.org/0000-0001-6814-2812>

Timothy H Murphy: <https://orcid.org/0000-0002-0093-4490>

References

- Abbasi Aamir**, Lassagne Henri, Estebanez Luc, Goueytes Dorian, Shulz Daniel E., Ego-Stengel Valerie (2023) Brain-Machine Interface Learning Is Facilitated by Specific Patterning of Distributed Cortical Feedback. *Science Advances* **9**:eadh1328
- Akam Thomas**, Lustig Andy, Rowland James M., Kapanaiah Sampath Kt, Esteve-Agraz Joan, Panniello Mariangela, Márquez Cristina, et al. (2022) Open-Source, Python-Based, Hardware and Software for Controlling Behavioural Neuroscience Experiments. *eLife* **11** <https://doi.org/10.7554/eLife.67846>
- Allen Cara B.**, Celikel Tansu, Feldman Daniel E. (2003) Long-Term Depression Induced by Sensory Deprivation during Cortical Map Plasticity in Vivo. *Nature Neuroscience* **6**:291-99
- Allen Christopher**, Mehler David M. A. (2019) Open Science Challenges, Benefits and Tips in Early Career and beyond. *PLoS Biology* **17**:e3000246
- Allen William E.**, Kauvar Isaac V., Chen Michael Z., Richman Ethan B., Yang Samuel J., Chan Ken, Gradinaru Viviana, Deverman Benjamin E., Luo Liqun, Deisseroth Karl (2017) Global Representations of Goal-Directed Behavior in Distinct Cell Types of Mouse Neocortex. *Neuron* **94**:891-907.
- Ching Shinung**, Liberman Max Y., Chemali Jessica J., Brandon Westover M., Kenny Jonathan D., Solt Ken, Purdon Patrick L., Brown Emery N. (2013) Real-Time Closed-Loop Control in a Rodent Model of Medically Induced Coma Using Burst Suppression. *Anesthesiology* **119**:848-60
- Chon Uree**, Vanselow Daniel J., Cheng Keith C., Kim Yongsoo (2019) Enhanced and Unified Anatomical Labeling for a Common Mouse Brain Atlas. *bioRxiv* <https://doi.org/10.1101/636175>

- Clancy Kelly B.**, Koralek Aaron C., Costa Rui M., Feldman Daniel E., Carmena Jose M. (2014) Volitional Modulation of Optically Recorded Calcium Signals during Neuroprosthetic Learning. *Nature Neuroscience* **17**:807-9
- Clancy Kelly B.**, Mrsic-Flogel Thomas D. (2021) The Sensory Representation of Causally Controlled Objects. *Neuron* **109**:677-89.
- Dana Hod**, Chen Tsai-Wen, Hu Amy, Shields Brenda C., Guo Caiying, Looger Loren L., Kim Douglas S., Svoboda Karel (2014) Thy1-GCaMP6 Transgenic Mice for Neuronal Population Imaging in Vivo. *PLoS One* **9**:e108697
- Dhillon Navjodh Singh**, Sutandi Agustinus, Vishwanath Manoj, Lim Miranda M., Cao Hung, Si Dong (2021) A Raspberry Pi-Based Traumatic Brain Injury Detection System for Single-Channel Electroencephalogram. *Sensors* **21** <https://doi.org/10.3390/s21082779>
- Fetz E. E** (1969) Operant Conditioning of Cortical Unit Activity. *Science* **163**:955-58
- Forys Brandon J.**, Xiao Dongsheng, Gupta Pankaj, Murphy Timothy H. (2020) Real-Time Selective Markerless Tracking of Forepaws of Head Fixed Mice Using Deep Neural Networks. *eNeuro* **7**:ENEURO.0096-20.2020
- Ganguly Karunesh**, Carmena Jose M. (2009) Emergence of a Stable Cortical Map for Neuroprosthetic Control. *PLoS Biology* **7**:e1000153
- Gilad Ariel**, Helmchen Fritjof (2020) Spatiotemporal Refinement of Signal Flow through Association Cortex during Learning. *Nature Communications* **11**:1744
- Girardeau Gabrielle**, Benchenane Karim, Wiener Sidney I., Buzsáki György, Zugaro Michaël B. (2009) Selective Suppression of Hippocampal Ripples Impairs Spatial Memory. *Nature Neuroscience* **12**:1222-23
- Han Yoon K.**, Köver Hania, Insanally Michele N., Semerdjian John H., Bao Shaowen (2007) Early Experience Impairs Perceptual Discrimination. *Nature Neuroscience* **10**:1191-97
- Huber D.**, Gutnisky D. A., Peron S., O'Connor D. H., Wiegert J. S., Tian L., Oertner T. G., Looger L. L., Svoboda K. (2012) Multiple Dynamic Representations in the Motor Cortex during Sensorimotor Learning. *Nature* **484**:473-78
- Kane Gary A.**, Lopes Gonçalo, Saunders Jonny L., Mathis Alexander, Mathis Mackenzie W. (2020) Real-Time, Low-Latency Closed-Loop Feedback Using Markerless Posture Tracking. *eLife* **9** <https://doi.org/10.7554/eLife.61909>
- Kim Tony Hyun**, Zhang Yanping, Lecoq Jérôme, Jung Juergen C., Li Jane, Zeng Hongkui, Niell Cristopher M., Schnitzer Mark J. (2016) Long-Term Optical Access to an Estimated One Million Neurons in the Live Mouse Cortex. *Cell Reports* **17**:3385-94
- Knudsen Eric B.**, Wallis Joni D. (2020) Closed-Loop Theta Stimulation in the Orbitofrontal Cortex Prevents Reward-Based Learning. *Neuron* **106**:537-47.
- Krakauer John W.**, Hadjiosif Alkis M., Xu Jing, Wong Aaron L., Haith Adrian M. (2019) Motor Learning. *Comprehensive Physiology* **9**:613-63
- Lee Candice**, Lavoie Andreanne, Liu Jiashu, Chen Simon X., Liu Bao-Hua (2020) Light Up the Brain: The Application of Optogenetics in Cell-Type Specific Dissection of Mouse Brain Circuits. *Frontiers in Neural Circuits* **14**:18
- Lim Diana H.**, Mohajerani Majid H., Ledue Jeffrey, Boyd Jamie, Chen Shangbin, Murphy Timothy H. (2012) In Vivo Large-Scale Cortical Mapping Using Channelrhodopsin-2 Stimulation in Transgenic Mice Reveals Asymmetric and Reciprocal Relationships between Cortical Areas. *Frontiers in Neural Circuits* **6**:11
- Lopes Gonçalo**, Monteiro Patricia (2021) New Open-Source Tools: Using Bonsai for Behavioral Tracking and Closed-Loop Experiments. *Frontiers in Behavioral Neuroscience* **15**:647640
- Luo Shengtao**, Zhou Xiaopeng, Zhou Hui, Li Tao, He Yutong, Chen Jiang-Fan, Zhang Liping (2024) "Volitional Modulation of Neuronal Activity in the External Globus Pallidus by Engagement of the Cortical-Basal Ganglia Circuit.". *The Journal of Physiology* <https://doi.org/10.1113/JP286046>

Makino Hiroshi, Hwang Eun Jung, Hedrick Nathan G., Komiyama Takaki (2016) Circuit Mechanisms of Sensorimotor Learning. *Neuron* **92**:705-21

Makino Hiroshi, Ren Chi, Liu Haixin, Kim An Na, Kondapaneni Neehar, Liu Xin, Kuzum Duygu, Komiyama Takaki (2017) Transformation of Cortex-Wide Emergent Properties during Motor Learning. *Neuron* **94**:880-90.

Mathis Alexander, Mamidanna Pranav, Cury Kevin M., Abe Taiga, Murthy Venkatesh N., Mathis Mackenzie Weygandt, Bethge Matthias (2018) DeepLabCut: Markerless Pose Estimation of User-Defined Body Parts with Deep Learning. *Nature Neuroscience* **21**:1281-89

Mathis Mackenzie Weygandt, Mathis Alexander, Uchida Naoshige (2017) Somatosensory Cortex Plays an Essential Role in Forelimb Motor Adaptation in Mice. *Neuron* **93**:1493-1503.

Ma Ying, Shaik Mohammed A., Kim Sharon H., Kozberg Mariel G., Thibodeaux David N., Zhao Hanzhi T., Yu Hang, Hillman Elizabeth M. C. (2016) Wide-Field Optical Mapping of Neural Activity and Brain Haemodynamics: Considerations and Novel Approaches. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* **371** <https://doi.org/10.1098/rstb.2015.0360>

Michelson Nicholas J., Bolaños Federico, Bolaños Luis A., Balbi Matilde, LeDue Jeffrey M., Murphy Timothy H. (2023) Meso-Py: Dual Brain Cortical Calcium Imaging in Mice during Head-Fixed Social Stimulus Presentation. *eNeuro* **10** <https://doi.org/10.1523/ENEURO.0096-23.2023>

Murphy Timothy H., Michelson Nicholas J., Boyd Jamie D., Fong Tony, Bolanos Luis A., Bierbrauer David, Siu Teri, et al. (2020) Automated Task Training and Longitudinal Monitoring of Mouse Mesoscale Cortical Circuits Using Home Cages. *eLife* **9**:May <https://doi.org/10.7554/eLife.55964>

Nakai Nobuhiro, Sato Masaaki, Yamashita Okito, Sekine Yukiko, Fu Xiaochen, Nakai Junichi, Zalesky Andrew, Takumi Toru (2023) Virtual Reality-Based Real-Time Imaging Reveals Abnormal Cortical Dynamics during Behavioral Transitions in a Mouse Model of Autism. *Cell Reports* **42**:112258

Neely Ryan M., Koralek Aaron C., Athalye Vivek R., Costa Rui M., Carmena Jose M. (2018) Volitional Modulation of Primary Visual Cortex Activity Requires the Basal Ganglia. *Neuron* **97**:1356-68.

Neely Ryan M., Piech David K., Santacruz Samantha R., Maharbiz Michel M., Carmena Jose M. (2018) Recent Advances in Neural Dust: Towards a Neural Interface Platform. *Current Opinion in Neurobiology* **50**:64-71

O'Connor Daniel H., Andrew Hires S., Guo Zengcai V., Li Nuo, Yu Jianing, Sun Qian-Quan, Huber Daniel, Svoboda Karel (2013) Neural Coding during Active Somatosensation Revealed Using Illusory Touch. *Nature Neuroscience* **16**:958-65

Paz Jeanne T., Davidson Thomas J., Frechette Eric S., Delord Bruno, Parada Isabel, Peng Kathy, Deisseroth Karl, Huguenard John R. (2013) Closed-Loop Optogenetic Control of Thalamus as a Tool for Interrupting Seizures after Cortical Injury. *Nature Neuroscience* **16**:64-70

Peters Andrew J., Chen Simon X., Komiyama Takaki (2014) Emergence of Reproducible Spatiotemporal Activity during Motor Learning. *Nature* **510**:263-67

Prsa Mario, Galiñanes Gregorio L., Huber Daniel (2017) Rapid Integration of Artificial Sensory Feedback during Operant Conditioning of Motor Cortex Neurons. *Neuron* **93**:929-39.

Rosin Boris, Slovik Maya, Mitelman Rea, Rivlin-Etzion Michal, Haber Suzanne N., Israel Zvi, Vaadia Eilon, Bergman Hagai (2011) Closed-Loop Deep Brain Stimulation Is Superior in Ameliorating Parkinsonism. *Neuron* **72**:370-84

Siegle Joshua H., López Aarón Cuevas, Patel Yogi A., Abramov Kirill, Ohayon Shay, Voigts Jakob (2017) Open Ephys: An Open-Source, Plugin-Based Platform for Multichannel Electrophysiology. *Journal of Neural Engineering* **14**:045003

Silasi Gergely, Boyd Jamie D., Bolanos Federico, LeDue Jeff M., Scott Stephen H., Murphy Timothy H. (2018) Individualized Tracking of Self-Directed Motor Learning in Group-Housed Mice Performing a Skilled Lever Positioning Task in the Home Cage. *Journal of Neurophysiology* **119**:337-46

- Silasi Gergely**, Xiao Dongsheng, Vanni Matthieu P., Chen Andrew C. N., Murphy Timothy H. (2016) Intact Skull Chronic Windows for Mesoscopic Wide-Field Imaging in Awake Mice. *Journal of Neuroscience Methods* **267**:141-49
- Srinivasan Shriya S.**, Maimon Benjamin E., Diaz Maurizio, Song Hyungeun, Herr Hugh M. (2018) Closed-Loop Functional Optogenetic Stimulation. *Nature Communications* **9**:5303
- Sun Guanghao**, Zeng Fei, McCartin Michael, Zhang Qiaosheng, Xu Helen, Liu Yaling, Chen Zhe Sage, Wang Jing (2022) Closed-Loop Stimulation Using a Multiregion Brain-Machine Interface Has Analgesic Effects in Rodents. *Science Translational Medicine* **14**:eabm5868
- Valley M. T.**, Moore M. G., Zhuang J., Mesa N., Castelli D., Sullivan D., Reimers M., Waters J. (2020) Separation of Hemodynamic Signals from GCaMP Fluorescence Measured with Wide-Field Imaging. *Journal of Neurophysiology* **123**:356-66
- Vanni Matthieu P.**, Murphy Timothy H. (2014) Mesoscale Transcranial Spontaneous Activity Mapping in GCaMP3 Transgenic Mice Reveals Extensive Reciprocal Connections between Areas of Somatomotor Cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* **34**:15931-46
- Wang Quanxin**, Ding Song-Lin, Li Yang, Royall Josh, Feng David, Lesnar Phil, Graddis Nile, et al. (2020) The Allen Mouse Brain Common Coordinate Framework: A 3D Reference Atlas. *Cell* **181**:936-53.
- Wekselblatt Joseph B.**, Flister Erik D., Piscopo Denise M., Niell Christopher M. (2016) Large-Scale Imaging of Cortical Dynamics during Sensory Perception and Behavior. *Journal of Neurophysiology* **115**:2852-66
- White Samantha R.**, Amarante Linda M., Kravitz Alexxai V., Laubach Mark (2019) The Future Is Open: Open-Source Tools for Behavioral Neuroscience Research. *eNeuro* **6**
<https://doi.org/10.1523/ENEURO.0223-19.2019>
- Widge Alik S.**, Moritz Chet T. (2014) Pre-Frontal Control of Closed-Loop Limbic Neurostimulation by Rodents Using a Brain-Computer Interface. *Journal of Neural Engineering* **11**:024001
- Wolpert Daniel M.**, Diedrichsen Jörn, Flanagan J. Randall (2011) Principles of Sensorimotor Learning. *Nature Reviews. Neuroscience* **12**:739-51
- Xiao Dongsheng**, Vanni Matthieu P., Mitelut Catalin C., Chan Allen W., LeDue Jeffrey M., Xie Yicheng, Chen Andrew C. N., Swindale Nicholas V., Murphy Timothy H. (2017) Mapping Cortical Mesoscopic Networks of Single Spiking Cortical or Sub-Cortical Neurons. *eLife* **6**:February
<https://doi.org/10.7554/elife.19976>
- Yu Che-Hang**, Stirman Jeffrey N., Yu Yiyi, Hira Riichiro, Smith Spencer L. (2021) Diesel2p Mesoscope with Dual Independent Scan Engines for Flexible Capture of Dynamics in Distributed Neural Circuitry. *Nature Communications* **12**:6639
- Kumar Pankaj Gupta**, Timothy Murphy (2024) Real-Time Closed-Loop Feedback System For Mouse Mesoscale Cortical Signal And Movement Control: CLoPy. FRDR.
<https://doi.org/10.20383/103.01152>

Peer reviews

Reviewer #1 (Public review):

Summary:

The authors provide a resource to the systems neuroscience community by offering their Python-based CLoPy platform for closed-loop feedback training. In addition to using neural feedback, as is common in these experiments, they include a capability to use real-time movement extracted from DeepLabCut as the control signal. The methods and repository are detailed for those who wish to use this resource. Furthermore, they demonstrate the efficacy of their system through a series of mesoscale calcium imaging experiments. These experiments use a large number of cortical regions for the control signal in the neural

feedback setup, while the movement feedback experiments are analyzed more extensively. The revised preprint has improved substantially upon the previous submission.

Strengths:

The primary strength of the paper is the availability of their CLoPy platform. Currently, most closed-loop operant conditioning experiments are custom built by each lab, and carry a relatively large startup cost to get running. This platform lowers the barrier to entry for closed-loop operant conditioning experiments, in addition to making the experiments more accessible to those with less technical expertise.

Another strength of the paper is the use of many different cortical regions as control signals for the neurofeedback experiments. Rodent operant conditioning experiments typically record from the motor cortex, and maybe one other region. Here, the authors demonstrate that mice can volitionally control many different cortical regions not limited to those previously studied, recording across many regions in the same experiment. This demonstrates the relative flexibility of modulating neural dynamics, including in non-motor regions.

Finally, adapting the closed-loop platform to use real-time movement as a control signal is a nice addition. Incorporating movement kinematics into operant conditioning experiments has been a challenge due to the increased technical difficulties of extracting real-time kinematic data from video data at a latency where it can be used as a control signal for operant conditioning. In this paper, they demonstrate that the mice can learn the task using their forelimb position, at a rate that is quicker than the neurofeedback experiments.

Weaknesses:

Many of the original weaknesses have been addressed in the revised preprint.

While the dataset contains an impressive amount of animals and cortical regions for the neurofeedback experiment, my excitement for these experiments is tempered by the relative incompleteness of the dataset.

Additionally, adoption of the platform may be hindered by the absence of a tutorial on how to run a session.

<https://doi.org/10.7554/eLife.105070.2.sa3>

Reviewer #2 (Public review):

Summary:

In this work, Gupta & Murphy present several parallel efforts. On one side, they present the hardware and software they use to build a head-fixed mouse experimental setup that they use to track in "real-time" the calcium activity in one or two spots at the surface of the cortex. On the other side, they present another setup that they use to take advantage of the "real-time" version of DeepLabCut with their mice. The hardware and software that they used/develop is described at length, both in the article and in a companion GitHub repository. Next, they present experimental work that they have done with these two setups, training mice to max out a virtual cursor to obtain a reward, by taking advantage of auditory tone feedback that is provided to the mice as they modulate either (1) their local cortical calcium activity, or (2) their limb position.

Strengths:

This work illustrates the fact that thanks to readily available experimental building blocks, body movement and calcium imaging can be carried out using readily available components,

including imaging the brain using an incredibly cheap consumer electronics RGB camera (RGB Raspberry Pi Camera). It is a useful source of information for researchers that may be interested in building a similar setup, given the highly detailed overview of the system. Finally, it further confirms previous findings regarding the operant conditioning of the calcium dynamics at the surface of the cortex (Clancy et al. 2020) and suggests an alternative based on deepabcut to the motor tasks that aim to image the brain at the mesoscale during forelimb movements (Quarta et al. 2022).

Weaknesses:

This work covers 3 separate research endeavors: (1) The development of two separate setups, their corresponding software. (2) A study that is highly inspired from the Clancy et al. 2021 paper on the modulation of the local cortical activity measured through a mesoscale calcium imaging setup. (3) A study of the mesoscale dynamics of the cortex during forelimb movements learning. Sadly, the analyses of the physiological data appears incomplete, and more generally, the paper shows weaknesses regarding several points:

The behavioral setups that are presented are representative of the state of the art in the field of mesoscale imaging/head fixed behavior community, rather than a highly innovative design. Still, they definitely have value as a starting point for laboratories interested in implementing such approaches.

Throughout the paper, there are several statements that point out how important it is to carry out this work in a closed-loop setting with an auditory feedback, but sadly there is no "no feedback" control in cortical conditioning experiments, while there is a no-feedback condition in the forelimb movement study, which shows that learning of the task can be achieved in the absence of feedback.

The analysis of the closed-loop neuronal data behavior lacks controls. Increased performance can be achieved by modulating actively only one of the two ROIs, this is not really analyzed, while this finding which does not match previous reports (Clancy et al. 2020) would be important to further examine.

<https://doi.org/10.7554/eLife.105070.2.sa2>

Reviewer #3 (Public review):

Summary:

The study demonstrates the effectiveness of a cost-effective closed-loop feedback system for modulating brain activity and behavior in head-fixed mice. Authors have tested real-time closed-loop feedback system in head-fixed mice two types of graded feedback: 1) Closed-loop neurofeedback (CLNF), where feedback is derived from neuronal activity (calcium imaging), and 2) Closed-loop movement feedback (CLMF), where feedback is based on observed body movement. It is a python based opensource system, and the authors call it CLoPy. Authors also claim to provide all software, hardware schematics, and protocols to adapt it to various experimental scenarios. This system is capable and can be adapted for a wide use case scenarios.

Authors have shown that their system can control both positive (water drop) and negative reinforcement (buzzer-vibrator). This study also shows that using the closed-loop system, mice have shown to better performance, learnt arbitrary tasks and can adapt to changes in the rules as well. By integrating real-time feedback based on cortical GCaMP imaging and behavior tracking authors have provided strong evidence that such closed-loop systems can be instrumental in exploring the dynamic interplay between brain activity and behavior.

Strengths:

Simplicity of feedback systems design. Simplicity of implementation and potential adoption.

Weaknesses:

Long latencies, due to slow Ca^{2+} dynamics and slow imaging (15 FPS), may limit the application of the system.

<https://doi.org/10.7554/eLife.105070.2.sa1>

Author response

The following is the authors' response to the original reviews.

Public reviews:

Reviewer #1 (Public review):

Summary:

The authors provide a resource to the systems neuroscience community, by offering their Python-based CLoPy platform for closed-loop feedback training. In addition to using neural feedback, as is common in these experiments, they include a capability to use real-time movement extracted from DeepLabCut as the control signal. The methods and repository are detailed for those who wish to use this resource. Furthermore, they demonstrate the efficacy of their system through a series of mesoscale calcium imaging experiments. These experiments use a large number of cortical regions for the control signal in the neural feedback setup, while the movement feedback experiments are analyzed more extensively.

Strengths:

The primary strength of the paper is the availability of their CLoPy platform. Currently, most closed-loop operant conditioning experiments are custom built by each lab and carry a relatively large startup cost to get running. This platform lowers the barrier to entry for closed-loop operant conditioning experiments, in addition to making the experiments more accessible to those with less technical expertise.

Another strength of the paper is the use of many different cortical regions as control signals for the neurofeedback experiments. Rodent operant conditioning experiments typically record from the motor cortex and maybe one other region. Here, the authors demonstrate that mice can volitionally control many different cortical regions not limited to those previously studied, recording across many regions in the same experiment. This demonstrates the relative flexibility of modulating neural dynamics, including in non-motor regions.

Finally, adapting the closed-loop platform to use real-time movement as a control signal is a nice addition. Incorporating movement kinematics into operant conditioning experiments has been a challenge due to the increased technical difficulties of extracting real-time kinematic data from video data at a latency where it can be used as a control signal for operant conditioning. In this paper they demonstrate that the mice can learn the task using their forelimb position, at a rate that is quicker than the neurofeedback experiments.

Weaknesses:

There are several weaknesses in the paper that diminish the impact of its strengths. First, the value of the CLoPy platform is not clearly articulated to the systems neuroscience

community. Similarly, the resource could be better positioned within the context of the broader open-source neuroscience community. For an example of how to better frame this resource in these contexts, I recommend consulting the pyControl paper. Improving this framing will likely increase the accessibility and interest of this paper to a less technical neuroscience audience, for instance by highlighting the types of experimental questions CLoPy can enable.

We appreciate the editor's feedback regarding the clarity of the CLoPy platform's value and its positioning within the broader neuroscience community. We agree and understand the importance of effectively communicating the utility of CLoPy to both the systems neuroscience field and the wider open-source neuroscience community.

To address this, we have revised the introduction and discussion sections of the manuscript to more clearly articulate the unique contributions of the CLoPy platform. Specifically:

(1) We have emphasized how CLoPy can address experimental questions in systems neuroscience by highlighting its ability to enable real-time closed-loop experiments, such as investigating neural dynamics during behavior or studying adaptive cortical reorganization after injury. These examples are aimed at demonstrating its practical utility to the neuroscience audience.

(2) We have positioned CLoPy within the broader open-source neuroscience ecosystem, drawing comparisons to similar resources like pyControl. We describe how CLoPy complements existing tools by focusing on real-time optical feedback and integration with genetically encoded indicators, which are becoming increasingly popular in systems neuroscience. We also emphasize its modularity and ease of adoption in experimental settings with limited resources.

(3) To make the manuscript more accessible to a less technically inclined audience, we have restructured certain sections to focus on the types of experiments CLoPy enables, rather than the technical details of the implementation.

We have consulted the pyControl paper, as suggested, and have used it as a reference point to improve the framing of our resource. We believe these changes will increase the accessibility and appeal of the paper to a broader neuroscience audience.

While the dataset contains an impressive amount of animals and cortical regions for the neurofeedback experiment, and an analysis of the movement-feedback experiments, my excitement for these experiments is tempered by the relative incompleteness of the dataset, as well as its description and analysis in the text. For instance, in the neurofeedback experiment, many of these regions only have data from a single mouse, limiting the conclusions that can be drawn. Additionally, there is a lack of reporting of the quantitative results in the text of the document, which is needed to better understand the degree of the results. Finally, the writing of the results section could use some work, as it currently reads more like a methods section.

Thank you for your thoughtful and constructive feedback on our manuscript. We appreciate the time and effort you took to review our work and provide detailed suggestions for improvement. Below, we address the key points raised in your review:

(1) Dataset Completeness: We acknowledge that some of the neurofeedback experiments include data from only a single mouse for some cortical regions while for some cortical regions, there are several animals. This was due to practical constraints during the study, and we understand the limitations this poses for drawing broad conclusions. We felt it was still important to include these data sets with smaller sample sizes as they might be useful for others pursuing this direction in the future. To address this, we have revised the text to explicitly acknowledge these limitations and clarify that the results for some regions are

exploratory in nature. We believe our flexible tool will provide a means for our lab and others include more animals representing additional cortical regions in future studies. Importantly, we have included all raw and processed data as well as code for future analysis.

(2) Quantitative Results: We recognize the importance of reporting quantitative results in the text for better clarity and interpretation. In response, we have added more detailed description of the quantitative findings from both the neurofeedback and movement-feedback experiments. This will include effect sizes, statistical measures, and key numerical results to provide a clearer understanding of the degree and significance of the observed effects.

(3) Results Section Writing: We appreciate your observation that parts of the results section read more like a methods section. To improve clarity and focus, we have restructured the results section to present the findings in a more concise and interpretative manner, while moving overly detailed descriptions of experimental procedures to the methods section.

Suggestions for improved or additional experiments, data or analyses:

Not necessary for this paper, but it would be interesting to see if the CLNF group could learn without auditory feedback.

This is a great suggestion and certainly something that could be done in the future.

There are no quantitative results in the results section. I would add important results to help the reader better interpret the data. For example, in: "Our results indicated that both training paradigms were able to lead mice to obtain a significantly larger number of rewards over time," You could show a number, with an appropriate comparison or statistical test, to demonstrate that learning was observed.

Thank you for pointing this out. We have mentioned quantification values in the results now, along with being mentioned in the figure legends, and we are quoting it in following sentences. "A $\Delta F/F_0$ threshold value was calculated from a baseline session on day 0 that would have allowed 25% performance. Starting from this basal performance of around 25% on day 1, mice (CLNF No-rule-change, $N=23$, $n=60$ and CLNF Rule-change, $N=17$, $n=60$) were able to discover the task rule and perform above 80% over ten days of training (Figure 4A, RM ANOVA $p=2.83e-5$), and Rule-change mice even learned a change in ROIs or rule reversal (Figure 4A, RM ANOVA $p=8.3e-10$, Table 5 for different rule changes). There were no significant differences between male and female mice (Supplementary Figure 3A)."

For: "Performing this analysis indicated that the Raspberry Pi system could provide reliable graded feedback within ~ 63 {plus minus} 15 ms for CLNF experiments." The LED test shows the sending of the signal, but the actual delay for the audio generation might be longer. This is also longer than the 50 ms mentioned in the abstract.

We appreciate the reviewer's insightful comment. The latency reported (~ 63 ms) was measured using the LED test, which captures the time from signal detection to output triggering on the Raspberry Pi GPIO. We agree that the total delay for auditory feedback generation could include an additional latency component related to the digital-to-analog conversion and speaker response. In our setup, we employ a fast Audiostream library written in C to generate the audio signal and expect the delay contribution to be negligible compared to the GPIO latency. Though we did not do this, it can be confirmed by an oscilloscope-based pilot measurement (for additional delay calculation). We have updated the manuscript to clarify that the 63 ± 15 ms value reflects the GPIO-triggered output latency, and we have revised the abstract to accurately state the delay as " ~ 63 ms" rather than 50 ms. This ensures consistency and avoids underestimation of the latency. We have corrected the LED latency for CLNF and CLMF experiments in the abstract as well.

It could be helpful to visualize an individual trial for each experiment type, for instance how the audio frequency changes as movement speed / calcium activity changes.

We have added Supplementary Figure 8 that contains this data where you can see the target cortical activity trace, target paw speed, rewards, along with the audio frequency generated.

The sample sizes are small ($n=1$) for a few groups. I am excited by the variety of regions recorded, so it could be beneficial for the authors to collect a few more animals to beef up the sample sizes.

We've acknowledged that some of the sample sizes are small. Importantly, we have included raw and processed data as well as code for future analysis. We felt it was still important to still include these data sets with smaller sample sizes as they might be useful for others pursuing this direction in the future.

I am curious as to why 60 trials sessions were used. Was it mostly for the convenience of a 30 min session, or were the animals getting satiated? If the former, would learning have occurred more rapidly with longer sessions?

This is a great observation and the answer is it was mostly due to logistical reasons. We tried to not keep animals headfixed for more than 45 minutes in each session as they become less engaged with long duration headfixed sessions. After headfixing them, it takes about 15 minutes to get the experiment going and therefore 30 - 40 minutes long recorded sessions seemed appropriate before they stop being engaged or before they get satiated in the task. We provided supplemental water after the sessions and we observed that they consumed water after the sessions so they were not fully satiated during the sessions even when they performed well in the task and got maximum rewards. We also had inter-trial rest periods of 10s that elongated the session duration. We think it would be interesting to explore the relationship between session duration(number of trials) and task learning progression over the days in a separate study.

Figure 4E is interesting, it seems like the changes in the distribution of ΔF was in both positive and negative directions, instead of just positive. I'd be curious as to the author's thoughts as to why this is the case. Relatedly, I don't see Figure 4E, and a few other subplots, mentioned in the text. As a general comment, I would address each subplot in the text.

We have split Figure 4 into two to keep the figures more readable. Previous Figure 4E-H are now Figure 5A-D in the revised manuscript. The online real-time CLNF sessions were using a moving window average to calculate $\Delta F/F_0$ and the figures were generated by averaging the whole recorded sessions. We have added text in Methods under "Online $\Delta F/F_0$ calculation" and "Offline $\Delta F/F_0$ calculation" sections making it clear about how we do our $\Delta F/F_0$ normalization based on average fluorescence over the entire session. Using this method of normalization does increase the baseline so that some peaks appear to be below zero. Additionally, it is unclear what strategy animals are employing to achieve the rule specific target activity. The task did not constrain them to have a specific strategy for cortical activation - they were rewarded as long as they crossed the threshold in target ROI(s). For example, in 2-ROI experiments, to increase ROI1-ROI2 target activity, they could increase activity of ROI1 relative to ROI2 or decreased activity of ROI1 relative to ROI1 - both would have led to a reward as long as the result crossed the threshold.

We have now addressed and added reference to the figures in the text in Results under "Mice can explore and learn an arbitrary task, rule, and target conditions" and "Mice can rapidly adapt to changes in the task rule" sections - thanks for pointing this out.

For: "In general, all ROIs assessed that encompassed sensory, pre-motor, and motor areas were capable of supporting increased reward rates over time," I would provide a visual summary showing the learning curves for the different types of regions.

We have rewritten this section to emphasize that these conclusions were based on pooled data from multiple regions of interest. The sample sizes for each type of region are different and some are missing. We believe it would be incomplete and not comparable to present this as a regular analysis since the sample sizes were not balanced. We would be happy to dive deeper into this and point to the raw and processed dataset if anyone would like to explore this further by GitHub or other queries.

Relatedly, I would further explain the fast vs slow learners, and if they mapped onto certain regions.

Mice were categorized into fast or slow learners based on the slope of learning over days (reward progression over the days) as shown in Supplementary Figure 3C,D. Our initial aim was not to probe cortical regions that led to fast vs slow learning but this was a grouping we did afterwards. Based on the analysis we did, the fast learners included the sensory (V1), somatosensory (BC, HL), and motor (M1, M2) areas, while the slow learners included the motor (M1, M2), and higher order (TR, RL) cortical areas. Testing all dorsal cortical areas would be prudent to establish their role in fast or slow learning and it is an interesting future direction.

Also I would make the labels for these plots (e.g. Supp Fig3) more intuitive, versus the acronyms currently used.

We have made more expressive labels and explained the acronyms below the Supplementary Figure 3.

The CLMF animals showed a decrease in latency across learning, what about the CLNF animals? There is currently no mention in the text or figures.

We have now incorporated the CLNF task latency data into both the Results text and Figure 4C. Briefly, task latency decreased as performance improved, increased following a rule change, and then decreased again as the animals relearned the task. The previous Figure 4C has been updated to Figure 4D, and the former Figure 4D has been moved to Supplementary Figure 4E.

Reviewer #2 (Public review):

Summary:

In this work, Gupta & Murphy present several parallel efforts. On one side, they present the hardware and software they use to build a head-fixed mouse experimental setup that they use to track in "real-time" the calcium activity in one or two spots at the surface of the cortex. On the other side, they present another setup that they use to take advantage of the "real-time" version of DeepLabCut with their mice. The hardware and software that they used/develop is described at length, both in the article and in a companion GitHub repository. Next, they present experimental work that they have done with these two setups, training mice to max out a virtual cursor to obtain a reward, by taking advantage of auditory tone feedback that is provided to the mice as they modulate either (1) their local cortical calcium activity, or (2) their limb position.

Strengths:

This work illustrates the fact that thanks to readily available experimental building blocks, body movement and calcium imaging can be carried using readily available

components, including imaging the brain using an incredibly cheap consumer electronics RGB camera (RGB Raspberry Pi Camera). It is a useful source of information for researchers that may be interested in building a similar setup, given the highly detailed overview of the system. Finally, it further confirms previous findings regarding the operant conditioning of the calcium dynamics at the surface of the cortex (Clancy et al. 2020) and suggests an alternative based on deep learning to the motor tasks that aim to image the brain at the mesoscale during forelimb movements (Quarta et al. 2022).

Weaknesses:

This work covers 3 separate research endeavors: (1) The development of two separate setups, their corresponding software. (2) A study that is highly inspired from the Clancy et al. 2020 paper on the modulation of the local cortical activity measured through a mesoscale calcium imaging setup. (3) A study of the mesoscale dynamics of the cortex during forelimb movements learning. Sadly, the analyses of the physiological data appears uncomplete, and more generally the paper tends to offer overstatements regarding several points:

In contrast to the introductory statements of the article, closed-loop physiology in rodents is a well-established research topic. Beyond auditory feedback, this includes optogenetic feedback (O'Connor et al. 2013, Abbasi et al. 2018, 2023), electrical feedback in hippocampus (Girardeau et al. 2009), and much more.

We have included and referenced these papers in our introduction section (quoted below) and rephrased the part where our previous text indicated there are fewer studies involving closed-loop physiology.

“Some related studies have demonstrated the feasibility of closed-loop feedback in rodents, including hippocampal electrical feedback to disrupt memory consolidation (Girardeau et al. 2009), optogenetic perturbations of somatosensory circuits during behavior (O'Connor et al. 2013), and more recent advances employing targeted optogenetic interventions to guide behavior (Abbasi et al. 2023).”

The behavioral setups that are presented are representative of the state of the art in the field of mesoscale imaging/head fixed behavior community, rather than a highly innovative design. In particular, the closed-loop latency that they achieve (>60 ms) may be perceived by the mice. This is in contrast with other available closed-loop setups.

We thank the reviewer for this thoughtful comment and fully agree that our closed-loop latency is larger than that achieved in some other contemporary setups. Our primary aim in presenting this work, however, is not to compete with the lowest possible latencies, but to provide an open-source, accessible, and flexible platform that can be readily adopted by a broad range of laboratories. By building on widely available and lower-cost components, our design lowers the barrier of entry for groups that wish to implement closed-loop imaging and behavioral experiments, while still achieving latencies well within the range that can support many biologically meaningful applications.

For example, our latency (~60 ms) remains compatible with experimental paradigms such as:

Motor learning and skill acquisition, where sensorimotor feedback on the scale of tens to hundreds of milliseconds is sufficient to modulate performance.

Operant conditioning and reward-based learning, in which reinforcement timing windows are typically broader and not critically dependent on sub-20 ms latencies.

Cortical state dependent modulation, where feedback linked to slower fluctuations in brain activity (hundreds of milliseconds to seconds) can provide valuable insight.

Studies of perception and decision-making, in which stimulus response associations often unfold on behavioral timescales longer than tens of milliseconds.

We believe that emphasizing openness, affordability, and flexibility will encourage widespread adoption and adaptation of our setup across laboratories with different research foci. In this way, our contribution complements rather than competes with ultra-low-latency closed-loop systems, providing a practical option for diverse experimental needs.

Through the paper, there are several statements that point out how important it is to carry out this work in a closed-loop setting with an auditory feedback, but sadly there is no "no feedback" control in cortical conditioning experiments, while there is a no-feedback condition in the forelimb movement study, which shows that learning of the task can be achieved in the absence of feedback.

We fully agree that such a control would provide valuable insight into the contribution of feedback to learning in the CLNF paradigm. In designing our initial experiments, we envisioned multiple potential control conditions, including No-feedback and Random-feedback. However, our first and primary objective was to establish whether mice could indeed learn to modulate cortical ROI activation through auditory feedback, and to further investigate this across multiple cortical regions. For this reason, we focused on implementing the CLNF paradigm directly, without the inclusion of these additional control groups. To broaden the applicability of the system, we subsequently adapted the platform to the CLMF experiments, where we did incorporate a No-feedback group. These results, as the reviewer notes, strengthen the evidence for the role of feedback in shaping task performance. We agree that the inclusion of a No-feedback control group in the CLNF paradigm will be crucial in future studies to further dissect the specific contribution of feedback to cortical conditioning.

The analysis of the closed-loop neuronal data behavior lacks controls. Increased performance can be achieved by modulating actively only one of the two ROIs, this is not clearly analyzed (for instance looking at the timing of the calcium signal modulation across the two ROIs. It seems that overall ROIs1 and 2 covariate, in contrast to Clancy et al. 2020. How can this be explained?

We agree that the possibility of increased performance being driven by modulation of a single ROI is an important consideration. Our study indeed began with 1-ROI closed-loop experiments. In those early experiments, while we did observe animals improving performance across days, we realized that daily variability in ongoing cortical GCaMP activity could lead to fluctuations in threshold-crossing events. The 2-ROI design was subsequently introduced to reduce this variability, as the target activity was defined as the relative activity between the two ROIs (e.g., ROI1 – ROI2). This approach offered a more stable signal by normalizing ongoing fluctuations. In our analysis of the early 2-ROI experiments, we observed that animals adopted diverging strategies to achieve threshold crossings. Specifically, some animals increased activity in ROI1 relative to ROI2, while others decreased activity in ROI2 to accomplish the same effect. Once discovered, each animal consistently adhered to its chosen strategy throughout subsequent training sessions. This was an early and intriguing observation, but as the experiments were not originally designed to systematically test this effect, we limited our presentation to the analysis of a small number of animals (shown in Figure 11). We have added details about this observation in our Results section as well, quoted below-

“In the 2-ROI experiment where the task rule required “ROI1 - ROI2” activity to cross a threshold for reward delivery, mice displayed divergent strategies. Some animals predominantly increased ROI1 activity, whereas others reduced ROI2 activity, both approaches leading to successful threshold crossing (Figure 11)”.

We hope this clarifies how the use of two ROIs helps explain the apparent covariation of the signals, and why some divergence from the observations of Clancy et al. (2020) may be expected.

Reviewer #3 (Public review):

Summary:

The study demonstrates the effectiveness of a cost-effective closed-loop feedback system for modulating brain activity and behavior in head-fixed mice. Authors have tested real-time closed-loop feedback system in head-fixed mice two types of graded feedback: 1) Closed-loop neurofeedback (CLNF), where feedback is derived from neuronal activity (calcium imaging), and 2) Closed-loop movement feedback (CLMF), where feedback is based on observed body movement. It is a python based opensource system, and authors call it CLoPy. The authors also claim to provide all software, hardware schematics, and protocols to adapt it to various experimental scenarios. This system is capable and can be adapted for a wide use case scenario.

Authors have shown that their system can control both positive (water drop) and negative reinforcement (buzzer-vibrator). This study also shows that using the close loop system mice have shown better performance, learnt arbitrary task and can adapt to change in the rule as well. By integrating real-time feedback based on cortical GCaMP imaging and behavior tracking authors have provided strong evidence that such closed-loop systems can be instrumental in exploring the dynamic interplay between brain activity and behavior.

Strengths:

Simplicity of feedback systems designed. Simplicity of implementation and potential adoption.

Weaknesses:

Long latencies, due to slow Ca²⁺ dynamics and slow imaging (15 FPS), may limit the application of the system.

We appreciate the reviewer's comment and agree that latency is an important factor in our setup. The latency arises partly from the inherent slow kinetics of calcium signaling and GCaMP6s, and partly from the imaging rate of 15 FPS (every 66 ms). These limitations can be addressed in several ways: for example, using faster calcium indicators such as GCaMP8f, or adapting the system to electrophysiological signals, which would require additional processing capacity. In our implementation, image acquisition was fixed at 15 FPS to enable real-time frame processing (256 × 256 resolution) on Raspberry Pi 4B devices. With newer hardware, such as the Raspberry Pi 5, substantially higher acquisition and processing rates are feasible (although we have not yet benchmarked this extensively). More powerful platforms such as Nvidia Jetson or conventional PCs would further support much faster data acquisition and processing.

Major comments:

(1) Page 5 paragraph 1: "We tested our CLNF system on Raspberry Pi for its compactness, general-purpose input/output (GPIO) programmability, and wide community support, while the CLMF system was tested on an Nvidia Jetson GPU device." Can these programs and hardware be integrated with windows-based system and a microcontroller (Arduino/Tency). As for the broad adaptability that's what a lot of labs would already have (please comment/discuss)?

While we tested our CLNF system on a Raspberry Pi (chosen for its compactness, GPIO programmability, and large user community) and our CLMF system on an Nvidia Jetson GPU device (to leverage real-time GPU-based inference), the underlying software is fully written in Python. This design choice makes the system broadly adaptable: it can be run on any device capable of executing Python scripts, including Windows-based PCs, Linux machines, and macOS systems. For hardware integration, we have confirmed that the framework works seamlessly with microcontrollers such as Arduino or Teensy, requiring only minor modifications to the main script to enable sending and receiving of GPIO signals through those boards. In fact, we are already using the same system in an in-house project on a Linux-based PC where an Arduino is connected to the computer to provide GPIO functionality. Furthermore, the system is not limited to Raspberry Pi or Arduino boards; it can be interfaced with any GPIO-capable devices, including those from Adafruit and other microcontroller platforms, depending on what is readily available in individual labs. Since many neuroscience and engineering laboratories already possess such hardware, we believe this design ensures broad accessibility and ease of integration across diverse experimental setups.

(2) Hardware Constraints: The reliance on Raspberry Pi and Nvidia Jetson (is expensive) for real-time processing could introduce latency issues (~63 ms for CLNF and ~67 ms for CLMF). This latency might limit precision for faster or more complex behaviors, which authors should discuss in the discussion section.

In our system, we measured latencies of approximately ~63 ms for CLNF and ~67 ms for CLMF. While such latencies indeed limit applications requiring millisecond precision, such as fast whisker movements, saccades, or fine-reaching kinematics, we emphasize that many relevant behaviors, including postural adjustments, limb movements, locomotion, and sustained cortical state changes, occur on timescales that are well within the capture range of our system. Thus, our platform is appropriate for a range of mesoscale behavioral studies that probably needs to be discussed more. It is also important to note that these latencies are not solely dictated by hardware constraints. A significant component arises from the inherent biological dynamics of the calcium indicator (GCaMP6s) and calcium signaling itself, which introduce slower temporal kinetics independent of processing delays. Newer variants, such as GCaMP8f, offer faster response times and could further reduce effective biological latency in future implementations.

With respect to hardware, we acknowledge that Raspberry Pi provides a low-cost solution but contributes to modest computational delays, while Nvidia Jetson offers faster inference at higher cost. Our choice reflects a balance between accessibility, cost-effectiveness, and performance, making the system deployable in many laboratories. Importantly, the modular and open-source design means the pipeline can readily be adapted to higher-performance GPUs or integrated with electrophysiological recordings, which provide higher temporal resolution. Finally, we agree with the reviewer that the issue of latency highlights deeper and interesting questions regarding the temporal requirements of behavior classification. Specifically, how much data (in time) is required to reliably identify a behavior, and what is the minimum feedback delay necessary to alter neural or behavioral trajectories? These are critical questions for the design of future closed-loop systems and ones that our work helps frame.

We have added a slightly modified version of our response above in the discussion section under “Experimental applications and implications”.

(3) Neurofeedback Specificity: The task focuses on mesoscale imaging and ignores finer spatiotemporal details. Sub-second events might be significant in more nuanced behaviors. Can this be discussed in the discussion section?

This is a great point and we have added the following to the discussion section. “In the case of CLNF we have focused on regional cortical GCAMP signals that are relatively slow in kinetics. While such changes are well suited for transcranial mesoscale imaging assessment, it is possible that cellular 2-photon imaging (Yu et al. 2021) or preparations that employ cleared crystal skulls (Kim et al. 2016) could resolve more localized and higher frequency kinetic signatures.”

(4) The activity over 6s is being averaged to determine if the threshold is being crossed before the reward is delivered. This is a rather long duration of time during which the mice may be exhibiting stereotyped behaviors that may result in the changes in DFF that are being observed. It would be interesting for the authors to compare (if data is available) the behavior of the mice in trials where they successfully crossed the threshold for reward delivery and in those trials where the threshold was not breached. How is this different from spontaneous behavior and behaviors exhibited when they are performing the test with CLNF?

We would like to emphasize that we are not directly averaging activity over 6 s to compare against the reward threshold. Instead, the preceding 6 s of activity is used solely to compute a dynamic baseline for $\Delta F/F_0$ ($\Delta F/F_0 = (F - F_0)/F_0$). Here, F_0 is calculated as the mean fluorescence intensity over the prior 6 s window and is updated continuously throughout the session. This baseline is then subtracted from the instantaneous fluorescence signal to detect relative changes in activity. The reward threshold is therefore evaluated against these baseline-corrected $\Delta F/F_0$ values at the current time point, not against an average over 6 s. This moving-window baseline correction is a standard approach in calcium imaging analyses, as it helps control for slow drifts in signal intensity, bleaching effects, or ongoing fluctuations unrelated to the behavior of interest. Thus, the 6-s window is not introducing a temporal lag in reward assignment but is instead providing a reference to detect rapid increases in cortical activity. We have added the term dynamic baseline to the Methods to clarify.

Recommendations for the authors

Reviewer #1 (Recommendations for the authors):

Additional suggestions for improved or additional experiments, data or analyses.

For: "Looking closely at their reward rate on day 5 (day of rule change), they had a higher reward rate in the second half of the session as compared to the first half, indicating they were adapting to the rule change within one session." It would be helpful to see this data, and would be good to see within-session learning on the rule change day

Thank you for pointing this out. We had missed referencing the figure in the text, and have now added a citation to Supplementary Figure 4A, which shows the cumulative rewards for each day of training. As seen in the plot for day 5, the cumulative rewards are comparable to those on day 1, with most rewards occurring during the second half of the session.

For: "These results suggest that motor learning led to less cortical activation across multiple regions, which may reflect more efficient processing of movement-related activity," it could also be the case that the behaviour became more stereotyped over learning, which would lead to more concentrated, correlated activity. To test this, it would be good to look at the limb variability across sessions. Similarly, if it is movement-related, there should be good decoding of limb kinematics.

Indeed, we observed that behavior became more stereotyped over the course of learning, as shown in Supplementary Figure 4C, 4D. One plausible explanation for the reduction in cortical activation across multiple regions is that behavior itself became more stereotyped, a

possibility we have explored in the manuscript. Specifically, forelimb movements during the trial became increasingly correlated as mice improved on the task, particularly in the groups that received auditory feedback (Rule-change and No-rule-change groups; Figure 8). As movements became more correlated, overall body movements during trials decreased and aligned more closely with the task rule (Figure 9D). This suggests that reduced cortical activity may in part reflect changes in behavior. Importantly, however, in the Rule-change group, we observed that on the day of the rule switch (day 5), when the target shifted from the left to the right forelimb, cortical activity increased bilaterally (Figure 9A–C). This finding highlights our central point: groups that received feedback (Rule-change and No-rule-change) were able to identify the task rule more effectively, and both their behavior and cortical activity became more specifically aligned with the rule compared to the No-feedback group. We agree with the reviewers that additional analyses along these lines would be valuable future directions. To facilitate this, we have included the movement data for readers who may wish to pursue further analyses, details can be found under “Data and code availability” in Methods section. However, given the limited sample sizes in our dataset and the need to keep the manuscript focused on the central message, we felt that including these additional analyses here would risk obscuring the main findings.

For: "We believe the decrease in $\Delta F/F_{0peak}$ is unlikely to be driven by changes in movement, as movement amplitudes did not decrease significantly during these periods (Figure 7D CLMF Rule-change)." I would formally compare the two conditions. This is an important control. Also, another way to see if the change in ΔF is related to movement would be to see if you can predict movement from the ΔF .

Figure 7D in the previous version is Figure 9D in the current revision of the manuscript. We've assessed this for the examples shown based on graphing the movement data, unfortunately there is not enough of that data to do a group analysis of movement magnitude. We would suggest that this would be an excellent future direction that would take advantage of the flexible open source nature of our tool.

Recommendations for improving the writing and presentation.

In the abstract there is no mention of the rationale for the project, or the resulting significance. I would modify this to increase readership by the behavioral neuroscience community. Similarly, the introduction also doesn't highlight the value of this resource for the field. Again, I think the pyControl paper does a good job of this. For readability, I would add more subheadings earlier in the results, to separate the different technical aspects of the system.

We have revised the introduction to include the rationale for the project, its potential implications, and its relevance for translational research. We have also framed the work within the broader context of the behavioral and systems neuroscience community. We greatly appreciate this suggestion, as we believe it enhances the clarity and accessibility of the manuscript for the community.

For: "While brain activity can be controlled through feedback, other variables such as movements have been less studied, in part because their analysis in real time is more challenging." I would highlight research that has studied the control of behavior through feedback, such as the Mathis paper where mice learn to pull a joystick to a virtual box, and adapt this motion to a force perturbation.

We have added a citation to the Mathis paper and describe this as an additional form of feedback. The text is quoted below:

“Opportunities also exist in extending real time pose classification (Forys et al. 2020; Kane et al. 2020) and movement perturbation (Mathis et al. 2017) to shape aspects of an animal's motor repertoire.”

Some of the results content would be better suited for the methods, one example: "A previous version of the CLNF system was found to have non-linear audio generation above 10 kHz, partly due to problems in the audio generation library and partly due to the consumer-grade speaker hardware we were employing. This was fixed by switching to the Audiostream (<https://github.com/kivy/audiostream>) library for audio generation and testing the speakers to make sure they could output the commanded frequencies"

This is now moved to the Methods section.

For: "There are reports of cortical plasticity during motor learning tasks, both at cellular and mesoscopic scales (17-19), supporting the idea that neural efficiency could improve with learning," not sure I agree with this, the studies on cortical plasticity are usually to show a neural basis for the learning observed, efficiency is separate from this.

We have modified this statement to remove the concept of efficiency "There are reports of cortical plasticity during motor learning tasks, both at cellular and mesoscopic scales (17-19)."

The paragraph that opens "Distinct task- and reward-related cortical dynamics" that describes the experiment should appear in the previous section, as the data is introduced there.

We have moved the mentioned paragraphs in the previous section where we presented the data and other experiment details. This makes the text more readable and contextual.

I would present the different ROI rules with better descriptors and visualization to improve the readability.

We have added Supplementary Figure 7, which provides visualizations of the ROIs across all task rules used in the CLNF experiments.

Minor corrections to the text and figures.

Figure 1 is a little crowded, combining the CLNF and CLMF experiments, I would turn this into a 2 panel figure, one for each, similar to how you did figure 2.

We have revised Figure 1 to include two panels, one for CLNF and one for CLMF. The colored components indicate elements specific to each setup, while the uncolored components represent elements shared between CLNF and CLMF. Relevant text in the manuscript is updated to refer to these figures.

For Figure 2, the organization of the CLMF section is not intuitive for the reader. I would reorder it so it has a similar flow as the CLNF experiment.

We have revised the figure by updating the layout of panel B (CLMF) to align with panel A (CLNF), thereby creating a more intuitive and consistent flow between the panels. We appreciate this helpful suggestion, which we believe has substantially improved the clarity of the figure. The corresponding text in the manuscript has also been updated to reflect these changes.

For Figure 3, highlight that C and E are examples. They also seem a little out of place, so they could even be removed.

We have now explicitly labeled Figures 3C and 3E as representative examples (figure legend and on figure itself). We believe including these panels provides helpful context for readers: Figure 3C illustrates how the ROIs align on the dorsal cortical brain map with segmented cortical regions, while Figure 3E shows example paw trajectories in three dimensions, allowing visualization of the movement patterns observed during the trials.

In the plots, I would add sample sizes, for instance, in CLNF learning curve in Figure 4A, how many animals are in each group?

We have labeled Figure 4 with number of animals used in CLNF (No-rule-change, N=23; Rule-change, N=17), and CLMF (Rule-change, N=8; No-rule-change, N=4; No-feedback, N=4).

Also, Figure 7 for example, which figures are single-sessions, versus across animals? For Figure 7c, what time bin is the data taken from?

We have clarified this now and mentioned it in all the figures. Figure 7 in the previous version is Figure 9 in the current updated manuscript. Figure 9A is from individual sessions on different days from the same mouse. Figure 9B is the group average reward centered $\Delta F/F_0$ activity in different cortical regions (Rule-change, N=8; No-rule-change, N=4; No-feedback, N=4). Figure 9C shows average $\Delta F/F_0$ peak values obtained within -1sec to +1sec centered around the reward point (N=8).

It says "punish" in Figure 3, but there is no punishment?

Yes, the task did not involve punishment. Each trial resulted in either a success, which is followed by a reward, or a failure, which is followed by a buzzer sound. To better reflect these outcomes, we have updated Figure 3 and replaced the labels "Reward" with "Success" and "Punish" with "Failure."

The regression on 5c doesn't look quite right, also this panel is not mentioned in the text.

The figure referred to by the reviewer as Figure 5 is now presented as Figure 6 in the revised manuscript. Regarding the reviewer's observation about the regression line in the left panel of Figure 5C, the apparent misalignment arises because the majority of the data points are densely clustered at the center of the scatter plot, where they overlap substantially. The regression line accurately reflects this concentration of overlapping data. To improve clarity, we have updated the figure and ensured that it is now appropriately referenced in the Results section.

Reviewer #2 (Recommendations for the authors):

(1) There would be many interesting observations and links between the peripheral and cortical studies if there was a body video available during the cortical study. Is there any such data available?

We agree that a detailed analysis of behavior during the CLNF task would be necessary to explore any behavior correlates with success in the task. Unfortunately, we do not have a sufficient video of the whole body to perform such an analysis.

(2) The text (p. 24) states: [intracortical GCAMP transients measured over days became more stereotyped in kinetics and were more correlated (to each other) as the task performance increased over the sessions (Figure 7E).] But I cannot find this quantification in the figures or text?

Figure 7 in the previous version of the manuscript now appears as Figure 9. In this figure, we present cortical activity across selected regions during trials, and in Figure 9E we highlight that this activity becomes more correlated. Since we did not formally quantify variability, we have removed the previous claim that the activity became stereotyped and revised the text in the updated manuscript accordingly.

Typos:

10-serest c (page 13)

Inverted color codes in figure 4E vs F

Reviewer #3 (Recommendations for the authors):

We have mostly attempted to limit the feedback to suggestions and posed a few questions that might be interesting to explore given the dataset the authors have collected.

Comments:

In close loop systems the latency is primary concern, and authors have successfully tested the latency of the system (Delay): from detection of an event to the reaction time was less than 67ms.

We have commented on the issues and limitations caused by latency, and potential future directions to overcome these challenges in responses to some of the previous comments.

Additional major comments:

"In general, all ROIs assessed that encompassed sensory, pre-motor, and motor areas were capable of supporting increased reward rates over time (Figure 4A, Animation 1)." Fig 4A is merely showing change in task performance over time and does not have information regarding the changes observed specific to CLNF for each ROI.

We acknowledge that the sample size for individual ROI rules was not sufficient for meaningful comparisons. To address this limitation, we pooled the data across all the rules tested. The manuscript includes a detailed list of the rules along with their corresponding sample sizes for transparency.

A $\Delta F/F_0$ threshold value was calculated from a baseline session on day 0 that would have allowed 25% performance. Starting from this basal performance of around 25% on day 1, mice (CLNF No-rule-change, $n=28$ and CLNF Rule-change, $n=13$). It is unclear what the replicates here are. Trials or mice? The corresponding Figure legend has a much smaller n value.

Thank you for pointing this out. We realized that we had not indicated the sample replicates in the figure, and the use of n instead of N for the number of animals may have been misleading. We have now corrected the notation and clarified this information in the figure to resolve the discrepancy.

What were the replicates for each ROI pairs evaluated?

Each ROI rule and number of mice and trials are listed in Table 5 and Table 6.

Our analysis revealed that certain ROI rules (see description in methods) lead to a greater increase in success rate over time than others (Supplementary Figure 3D). The Supplementary figures 3C and 3D are blurry and could use higher resolution images.

We have increased the font size of the text that was previously difficult to read and re-exported the figure at a higher resolution (300 DPI). We believe these changes will resolve the issue.

Also, It will help the reader is a visual representation of the ROI pairs are provided, instead of the text view. One interesting question is whether there are anatomical biases to fast vs slow learning pairs (Directionality - anterior/posterior, distance between the selected ROIs etc). This could be interesting to tease apart.

We have added Supplementary Figure 7, which provides visualizations of the ROIs across all task rules used in the CLNF experiments. While a detailed investigation of the anatomical basis of fast versus slow learning cortical ROIs is beyond the scope of the present study, we agree that this represents an exciting future direction for further research.

How distant should the ROIs be to achieve increased task performance?

We appreciate this insightful question. We did not specifically test this scenario. In our study, we selected 0.3×0.3 mm ROIs centered on the standard AIBS mouse brain atlas (CCF). At this resolution, ROIs do not overlap, regardless of their placement in a two-ROI experiment. Furthermore, because our threshold calculations are based on baseline recordings, we expect the system would function for any combination of ROI placements. Nonetheless, exploring this systematically would be an interesting avenue for future experiments.

Figures:

I would leave out some of the methodological details such as the protocol for water restriction (Fig. 3) out of the legend. This will help with readability.

We have removed some of the methodological details, including those mentioned above, from the legend of Figure 3 in the updated manuscript.

Fig 1 and Fig 2: In my opinion, It would be easier for the reader if the current Fig. 2, which provides a high level description of CLNF and CLBF is presented as Fig. 1. The current Fig. 1, goes into a lot of methodological implementation details, and also includes a lot of programming jargon that is being introduced early in the paper that is hard to digest early on in the paper's narrative.

Thank you for the suggestion. In the new manuscript, Figure 1 and Figure 2 have been swapped.

Higher-resolution images/ plots are needed in many instances. Unsure if this is the pdf compression done by the manuscript portal that is causing this.

All figures were prepared in vector graphics format using the open-source software Inkscape. For this manuscript, we exported the images at 300 DPI, which is generally sufficient for publication-quality documents. The submission portal may apply additional processing, which could have resulted in a reduction in image quality. We will carefully review the final submission files and ensure that all figures are clear and of high quality.

The authors repeatedly show ROI specific analysis M1_L, F1_R etc. It will be helpful to provide a key, even if redundant in all figures to help the reader.

We have now included keys to all such abbreviations in all the figures.

There are also instances of editorialization and interpretation e.g., "Surprisingly, the "Rule-change" mice were able to discover the change in rule and started performing above 70% within a day of the rule change, on day 6" that would be more appropriate in the main body of the paper.

Thank you for pointing this out in the figure legend, and we have removed it now since we already discussed this in the Results.

Minor comments

(1) The description of Figure 1 is hard to follow and can be described better based on how the information is processed and executed in the system from source to processing and back. Using separated colors (instead of shaded of grey) for the neuro feedback and

movement feedback would help as well. Common components could have a different color. The specification like the description of the config file should come later.

Figure 1 in the previous version is Figure 2 in the updated version. We have taken suggestions from other reviewers and made the figure easier to understand and split it into two panels with color coding Green for CLNF, Pink for CLMF specific parts while common shared parts are left without any color.

(2) Page 20 last paragraph:

Authors are neglecting that the rule change is done one day prior and the results that you see in the second half on the 6th day are not just because of the first half of the 6th day instead combined training on the 5th day (rule change) and then the first half of the 6th day. Rephrasing this observation is essential.

We have revised the text for clarity to indicate that the performance increase observed on day 6 is not necessarily attributable to training on that day. In fact, we noted and mentioned that mice began to perform the task better during the second half of the session on day 5 itself.

(3) The method section description of the CLMF setup (Page no 39 first paragraph) is more detailed, a diagram of this setup would make it easy to follow and a better read.

We have made changes to the CLMF setup (Figure 1B) and CLMF schematic (Figure 2B) to make it easier to understand parts of the setup and flow of control.

<https://doi.org/10.7554/eLife.105070.2.sa0>