New skill acquisition does not perturb existing function in sensory cortex

Brian B. Jeon[1-3], Steven M. Chase[1-3], and Sandra J. Kuhlman[1-4]
[1]Department of Biomedical Engineering, [2]Center for the Neural Basis of Cognition, [3]Neuroscience Institute, [4]Department of Biological Sciences, Carnegie Mellon University
Corresponding author: Sandra J. Kuhlman

Abstract

The ability to learn and practice new skills applies to the computation of abstract properties such as geometric form and neuroprosthetic control, as well as physical actions. We trained adult mice to successfully use an optical brain-computer interface (BCI) device so that we could assess the extent to which abstract skill learning disrupts previously acquired cortical function. We found that animals used multiple neural strategies to gain volitional control of the BCI device, and continued to do so after additional training. Despite significant engagement of local neural activity during BCI performance, tuning properties and stimulus encoding were not perturbed by new skill acquisition. These results indicate that flexible selection of neural strategies during goal-directed practice may facilitate the integration of new skills with existing function.

Introduction

Successful integration of new skills into an existing network requires balancing the maintenance of perception and action with the acquisition of new function. New skill learning, whether it is perceptual, motor, or abstract, involves changes in cellular physiology that are distributed across the brain[1–4]. Locally, up to 50% of the neurons within a circuit are recruited during task-specific behaviors, and it is observed that synaptic plasticity is pervasive among the activated neurons, particularly in the early phases of training[3,5]. The extent to which these widespread adaptive changes disrupt or otherwise impact existing circuit function is not known. Although addressing this issue is fundamental to understanding the biological constraints on learning, it is challenging to investigate because it requires identifying the neurons responsible for improved performance and monitoring these same neurons longitudinally throughout learning, as well as monitoring existing network function before and after new skill acquisition. We met these challenges by combining 2-photon calcium imaging with a brain computer interface (BCI) behavioral paradigm. Mice were trained to perform a *de novo* task wherein they earned a reward by modulating the activity of a selected set of sensory neurons in a user-defined pattern. Stimulus tuning and discriminability were quantified before and after skill acquisition to assess maintenance of visual function. The use of a genetically encoded sensor, which was expressed specifically in excitatory neurons, allowed the same neurons to be tracked throughout training.

We focused our study on the primary visual cortex (V1) because the functional output of this circuit is well-characterized, and therefore training-induced alterations can be readily detected. For example in adult animals, enhancing eye-specific input by transiently closing one eye destabilizes orientation preference in single neurons as well as signal correlation between neurons[6]. Modification is not limited to deprivation paradigms. Sensory representations are persistently altered by reinforced behavioral training under conditions in which specific visual

cues are associated with salience[7–11], and learning impacts the representation of non-trained stimuli as well[12]. These studies raise the possibility that every time new associations are made, existing function is altered. On the other hand, recruitment of new neurons selective for task-relevant stimuli has been shown to be context-dependent in training paradigms that cue reward contingencies[7] or carry context-specific information[13]. These observations are an indication that new associations could be learned without perturbing existing function, in situations where the task goals are recognized by the subject to be distinct from previously learned behaviors and associations.

We found that *de novo* skill learning engaged the majority of V1 neurons in our imaging field of view. Despite this widespread engagement, the stability of neural tuning to visual stimuli was not perturbed. Furthermore, response amplitude of individual neurons during vision remained stable, and decoding analysis revealed that the estimated amount of information carried in V1 was unchanged. Thus, stimulus representation was robust to new skill learning. Our results directly demonstrate that a new skill can be integrated into existing networks without disrupting previously developed function.

Results

We chose to train animals on an abstract task using a BCI so that we could reinforce a pattern of neural activity that was distinct from neural patterns associated with salience in the animals' home-cage conditions. As such, training was a unique experience, different from previously learned relationships, and therefore allowed us to unambiguously test whether perturbation of existing function is inherent to sensory learning. It is well-established that adult mice can learn to control BCIs using neurons in primary sensory cortex as well as motor cortex, in a goal-oriented manner that relies on the same plasticity mechanisms that are engaged during perceptual and motor learning[2,4,14–18]. An additional advantage of using a BCI is that we were able to couple the device to neural activity such that multiple strategies could lead to target acquisition, thereby making it possible to assess the flexibility by which neural strategies are selected during practice, both within and across training sessions. Prior to BCI training, 2 baseline visual stimulation sessions were recorded. A third visual stimulation session was recorded after successful BCI learning.

To facilitate BCI learning, mice were first conditioned to associate a tone with a reward. Six mice were subjected to a Go/No-go pitch association task in which a 15 kHz pitch was associated with a water reward, and a 5 KHz pitch was associated with the absence of reward. Five of the six mice learned the task (**Fig. S1**) and proceeded to BCI training within 8 days of learning the pitch association**.** In the BCI paradigm, the activity of six randomly chosen V1 glutamatergic excitatory neurons was directly coupled to a one-dimensional auditory feedback cursor (**Fig. 1A**). The BCI task was performed in the dark. The same six neurons were used to drive the BCI throughout the experiment (**Fig. 1B**). Three of the six neurons were assigned a positive value such that increases in their activity drove the cursor closer to target threshold (direct positive neurons, DP 1-3), and three neurons were assigned a negative value such that decreases in their activity moved the cursor towards target threshold (direct negative neurons, DN 1-3). The z-scored activity of each of the six neurons was weighted equally. When the combined neural trajectory of the six direct neurons reached the target threshold, a water reward was available for collection. The animal had 10 seconds to reach the target, if the target

was not reached in that time period, the feedback tone was turned off, the system was re-armed, and another trial was primed after an inter-trial period (see Methods). Including DN neurons ensured that the animal could not solve the task by simply increasing the activity of all neurons in the region. During the course of BCI training, one mouse was removed due to a procedural issue which prevented cross-session alignment of direct neurons. The remaining four mice learned the task within 10 sessions. Performance, defined as the frequency of trials in which the target was reached, significantly increased across sessions (**Figs. 1C, S2**). Linear regression of performance was computed on an animal-by-animal basis (α=0.05). Using this method to assess learning, rather than relying on analysis that pools all animals[17], we were able to identify the training session in which each animal demonstrated volitional control of the device. This session is referred to as the 'learning point' (LP). We also found that training reduced the latency to reach target threshold (**Fig. 1D**).

Volitional control of BCI is associated with a change in excitability of V1 neurons

We found that after BCI learning, the majority of V1 neurons exhibited a change in excitability. Eighty-one percent of the neurons that were tracked on the first session and the learning point session, including both indirect and direct neurons pooled across all 4 animals, had a significant change in their level of activity at the time of success. The majority of these neurons, 78.8%, increased their excitability, and the remaining neurons decreased their excitability (**Fig. 2A,B**). These observations demonstrate that successful BCI training induced modifications within V1 such that the pattern of network activity produced at the time of success was distinct between the naive and trained state, similar to observations made in motor cortex[19].

In three of the four mice we examined whether V1 neurons gained sensitivity to the auditory pitches used in the BCI paradigm. V1 neurons can be modulated by auditory stimuli that co-occur with visual stimuli; however, in the absence of training, it is well established that auditory stimuli do not drive V1 neurons to spike threshold[20–25]. We found that after BCI learning, V1 neurons were not significantly modulated by any of the 6 auditory pitches used in the training paradigm (**Fig. S3**). Thus, it is unlikely that BCI performance was improved due to an increased sensitivity to the auditory pitches used here.

To further explore the basis for increased performance on the BCI task, the summed activity of the direct negative and direct positive neurons was compared before and after learning. At the time of success, in three of the four mice the summed activity of the direct positive neurons was significantly increased, visualized as a rightward shift of the cumulative probability distribution (**Fig. 2C**). In one mouse, the summed activity of the direct negative neurons significantly decreased. Thus, in all four mice, a change in direct neuron activity levels could account for improved performance.

BCI training did not enhance stimulus responses

BCI-induced plasticity within V1 could influence visual function in three ways. Vision could be enhanced, disrupted, or maintained. To distinguish between these possibilities, first we considered whether visual responses in the subset of indirect neurons that were similarly tuned to the direct neurons were enhanced. Pairs of neurons that share orientation preference are preferentially connected at the synaptic level and exhibit correlated activity[26–28], therefore it is possible that BCI training could induce Hebbian plasticity in indirect-direct neuron pairs.

Indeed, other paradigms in which user-defined *de novo* correlations are introduced, Hebbian plasticity is evoked[29]. In our case, Hebbian plasticity would be revealed as a strengthening of responsiveness to the preferred visual stimulus of the direct neuron. To assess whether this was the case, we examined the response reliability of the indirect neurons to the preferred stimulus of the direct neurons, before and after the BCI task was learned.

Tuning to randomly presented static grating stimuli was assessed by fitting deconvolved responses to a two-dimensional Gaussian function[30](**Fig. 3A,B**). The stimulus set consisted of 12 orientations and 15 spatial frequencies spanning a range of 0.02-0.30 cycles/°, resulting in a total of 180 stimuli. Out of the 292 neurons that were tracked in both visual stimulation episodes across the 4 mice that learned the BCI task, 107 were tuned to grating stimuli, and out of those tuned neurons, 39 had similar orientation tuning preference to a direct neuron (**Fig. 3C-F**). We found that learning the BCI task did not increase the reliability of these 39 indirect neurons (**Fig. 3D**), nor was the response amplitude increased after removing failure trials (**Fig. 3E**). In addition, we found that across the four mice, the fraction of neurons similarly tuned to direct neurons was not increased by BCI training; prior to BCI training, 34±13% (±S.E.M.) of tuned neurons were similarly tuned to a direct neuron, and after learning the BCI task 35±12% neurons were similarly tuned.

To further ascertain whether tuning to visual stimuli may have been enhanced in V1, we examined the discriminability of grating stimulus orientation at the preferred spatial frequency of the direct neurons using decoding methods. A k-nearest neighbor (KNN) classifier was constructed to decode orientation at the preferred spatial frequency for each of the direct neurons that had a clear stimulus preference. Orientation discriminability was not improved following BCI training (**Fig. 3F**).

Stimulus representation was not disrupted after learning

To determine whether BCI training perturbed stimulus representation, four key aspects of visual processing were assessed: stability of tuning, pairwise correlation of neural activity during vision, response amplitude, and discriminability of visual features using decoding methods. Previously, we and others have characterized the stability of orientation and spatial frequency tuning. Despite high trial-to-trial variability in response amplitude in awake conditions, orientation and spatial frequency are stable across days[6,30,31]. The majority of tuned V1 neurons shift their preferred orientation five degrees or less over the course of two weeks, and similarly, preferred spatial frequency changes less than 0.006 cycles/° over the same period[30]. Here we show that in the adult, this level of stability does not require maintenance of visual input. A separate cohort of 6 adult mice not used for BCI experiments were held in constant darkness for 8 days (**Fig. S4**). Stability of tuning in response to grating stimuli was assessed. We found that shifts in tuning parameters were not different in baseline and dark-exposure conditions. Thus, the 2-3 hours of darkness experienced during BCI training is unlikely to impact tuning stability.

With this in mind, we assessed the stability of tuning in response to grating stimuli before and after BCI learning. A total of three imaging episodes were collected. Two episodes were collected before BCI training to establish a baseline and the third episode was collected after the BCI learning point. Across the first two episodes (episode 1 and 2), 353 neurons were tracked, and across the last two episodes (episode 2 and 3), 292 neurons were tracked (**Fig.**

4A,B). The stability of orientation and spatial frequency preference, as well as bandwidth for both of these features was compared between baseline and BCI conditions for neurons that were well-fit by the two-dimensional Gaussian function, in others words tuned, and responded to at least one of the 180 stimulus presentations (**Figs. 4C-D,S5**). Episode 2 was collected within 4 days of the completion of the auditory Go/No-Go task prior to the initiation of BCI training, and episode 3 was collected after the learning point in all 4 mice. This design ensured that the number of days spanning the baseline and BCI condition were matched within subjects and accounted for individual differences in the learning rate of both the auditory pitch association task as well as the BCI task. The median difference in orientation preference of the 353 neurons imaged during the baseline condition, pooled across mice was 8.4±1.9° (±S.E.M. across animals). Considering that the time span was slightly longer than previously published work[30], this was expected. We found that the distribution of changes in orientation preference for individual neurons during the BCI condition was largely overlapping with that of the baseline condition. The median difference in orientation preference in the BCI condition was 5.7±0.9° (**Fig. 4C**). Analysis of the distributions on an animal-by-animal basis revealed that orientation tuning was not destabilized by BCI learning (**Fig. 4D**). This was the case for the other tuning features as well (**Fig. S5**).

Previous work demonstrated that for a given neuron, the stability of preferred orientation and preferred spatial frequency are independently regulated[30]. Therefore, although individual tuning parameters were stable during BCI acquisition, it is possible that a neuron could be characterized as unstable when considering all four parameters together. To address this possibility, we examined pairwise signal correlation among the neurons that were both tuned and responsive. Signal correlation is the correlation of the average responses to the presented stimuli between a pair of neurons, as such is a measurement that is independent of the Gaussian model of tuning and is non-parametric. To assess the stability of signal correlation, we computed the change in signal correlation between baseline and BCI conditions. The distribution of changes in signal correlation was not different between the two conditions (**Fig. 4E**). Both distributions were centered on zero; the standard deviation of the baseline condition was 0.103, and the standard deviation of the BCI condition was 0.111.

Next we considered whether BCI training may have altered the strength of responsiveness. Altered response strength could modify the readout of downstream target areas. To address whether BCI training altered the strength of response, we examined the changes in the average amplitude of the preferred stimulus. In this case, we considered all neurons that were responsive, including those that were not necessarily well-tuned to orientation or spatial frequency. We found that changes in the response amplitude over the course of the baseline and BCI conditions was largely overlapping (**Fig. 4F**). The median change in amplitude was 1.77 events/bin in the baseline condition, and 1.88 events/bin in the BCI condition. Thus, acquisition of BCI control did not induce changes in response strength of V1 neurons.

Finally, a remaining possibility is that BCI training could have modified the amount of visual information contained within V1. To address this possibility, we utilized decoding methods to estimate the amount of information encoded by a population of V1 neurons, including neurons not well-tuned for orientation or spatial frequency. We used a k-nearest neighbor (KNN) classifier to decode stimulus orientation at a spatial frequency of 0.04 cycles/º,

the preferred spatial frequency of the majority of V1 neurons[32]. We found that decoding accuracy was not different after BCI acquisition compared to before BCI training in the same mice (**Fig. 4G**). In addition, decoding accuracy after BCI acquisition was indistinguishable from that of 4 control mice that never experienced BCI training. These results indicate that the amount of information encoded by the V1 population was not altered by BCI learning.

Flexible selection of neural strategies was maintained throughout training

In our BCI design, there are three broad classes of possible neural trajectories that can bring the auditory cursor to target threshold (**Fig. 5A**). First, both the direct positive and the direct negative neurons contribute to bringing the cursor closer to target threshold. This trajectory requires equal amounts of modulation across all direct neurons. Second, only the direct positive neurons contribute to bringing the cursor to target threshold, and third, only the direct negative neurons move the cursor to the target threshold (see Methods for how the trajectory classes are defined). In the example mouse shown in Figure 5A, on the first training session we found that all three types of trajectories were used, although trials in which only the direct positive neurons contributed to target acquisition were the most frequent. Across all mice, the most frequent trajectory, regardless of training status, was trajectory type 2, trials in which only the direct positive neurons moved the cursor to target threshold (**Figs. 5A,S6A**).

Closer examination of how successful trajectories were generated revealed that during training, different strategies were used to generate a given trajectory type. We quantified this variation by categorizing neural activity into 8 strategies. The 8 strategies accounted for 94.7% of all successful trials across the 4 mice (**Fig. 5B**). Single-trial strategies were defined as follows: a single direct neuron contributed to 70% or more of the neural drive required to reach the target (accounting for 6 categories, because there are 6 direct neurons), or either the 3 direct positive or the 3 direct negative neurons coordinated to contribute to 70% or more of the neural drive required to reach the target (accounting for 2 categories). In a minority of cases the strategy did not fall into one of these categories, such trials are referred to as unclassified. In 3 of the 4 mice, the activity of a single direct positive neuron contributed to 70% or more of the neural drive at the time of success for the majority of the trials. In one mouse, coordination among direct positives was the dominate strategy employed.

Next we determined whether the variability of strategies used within a given session decreased after learning. Training was extended for 4-9 sessions after the learning point. Prior to the learning point, on average 4.05 ± 0.06 (±S.E.M., across mice) of the 8 strategies were employed at least once in a given session. After the learning point, on average 4.1 ± 0.08 strategies were employed. Thus within a given session, a similar number of strategies were used before and after learning. For the three mice in which a change in the excitability of direct positive neurons was associated with increased performance **(Fig. 2C,** mice #1-3), specific strategies using direct positive neurons were identified as being correlated with performance (**Fig. S6B**). However, successful trials after the learning point were not dominated by the strategy that was associated with higher performance. In other words, the strategy that was associated with higher performance did not increase in frequency with extended training after the learning point. In the case of mouse #1, strategy ‘DP3’ was correlated with performance, yet was not the dominant strategy used in the majority of the subsequent sessions after the learning point. Similarly, in the cases of mouse #2 and #3 in which the strategy using the

coordinated activity of the direct positive neurons, '∑DP', was correlated with performance, this strategy was used in less than 50% of the successful trials after the learning point. Thus, multiple strategies continued to be used throughout training, even after learning.

Discussion

Theoretically, learning new skills could result in changes to neural activity that minimally impact previously acquired behaviors and function[33,34]. For example, synaptic plasticity could be restricted to updating activity patterns that occupy the null space of skilled action plans[35]. However, experimental evidence establishing that new skills can be acquired without altering the neural activity that underlies previously developed computations and function is lacking. Furthermore, learning in artificial network simulations often leads to catastrophic forgetting[36,37], raising the possibility that networks cannot integrate new information without disrupting previously learned functions. Here we directly demonstrated that integration of a new skill does not perturb previously developed function. We were able to address this issue by designing a learning task that associated a user-defined pattern of activity with a rewarding outcome, in combination with tracking the activity of the same neurons across skill learning as well as assess the original functional properties of the network before and after acquisition of the new skill. The original function of the network was probed at the population level as well as well as at the level of individual neuron properties. Learning the new skill did not destabilize neuronal tuning properties nor did it impact the network's ability to discriminate stimulus orientation, assessed using decoding methods. The task that we designed required modulation of a few neurons, yet we observed that training induced wide-spread changes within our imaging field of view, including changes in neurons that were not directly coupled to the device. Despite this wide-spread plasticity, we found that the encoding of visual stimuli was not altered by the acquisition of a new skill.

Furthermore, our results establish that in adult mice, stability of stimulus representation does not require on-going visual experience. We found that the stability of tuning to grating stimuli was not affected by one week of continuous dark exposure. Our data indicate that temporary changes in tuning stability induced by monocular deprivation in adults[6,38] is due to an increased drive from open eye inputs, rather than a loss of activity from closed eye inputs. It will be of interest to determine whether stability of stimulus tuning is maintained in the absence of input for other sensory modalities as well.

Does the stability of previous function observed here apply to association networks outside of the primary sensory cortex? We chose to perform our experiments in primary visual cortex because responses to visual stimuli are well characterized and known to be stable across days[6,30,31], thus any perturbation of function would be straightforward to detect. Notably, vision itself is an acquired function that requires experience-dependent instruction[39–42]. During early postnatal development, the circuits that process visual stimuli are refined following Hebbian rules, instructed by the input statistics experienced by the individual. In the primary cortex, experience-dependent properties include orientation selectivity and orientation discriminability [43,44], properties directly measured in this study, as well as ocular dominance[45,46], spatial receptive field position of ipsilateral input[47], and sensitivity to complex surround[48]. Taken together with observations that acquisition of BCI control recruits the same plasticity

mechanisms required for classic skill learning[2,4,14–18], it is reasonable to expect that the findings of this study apply generally to cortical circuits engaged in reinforced learning.

Within a single session, we found that various strategies were engaged to reach the target threshold. This trial-to-trial variability in neural activity pattern was maintained throughout training, even after the skill was established. In motor systems, variability during practice generates motor exploration that is known to facilitate trial-and-error learning[49–52]. Furthermore, theoretical studies indicate that after a new skill or function is acquired, behaviors are made more robust to changes in environmental conditions by continuing to use multiple strategies to perform the task. Mechanistically, variability in neural activity patterns ensures that strong correlations do not prevent adaptation in response to small perturbations during performance[53–55]. Our results suggest that variability during training may have an additional benefit. Engaging different activity patterns to achieve the same goal could facilitate the integration of new skills with existing function such that the network can integrate new skills without disturbing previously acquired functionality. Including variability during skill execution may be a key principle to incorporating increasingly complex routines during goal-directed behavior.

Author contributions

BBJ and SJK designed the experiments; BBJ collected the data; BBJ, SMC, and SJK analyzed the data and wrote the manuscript.

Acknowledgements

We thank Alison Barth, Aaron Batista, and Emily Oby for useful discussion. Funded by: NIH R01EY024678 (SJK), The Curci Foundation (SJK and SMC), and Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/Interior Business Center (DoI/IBCContract D16PC00007 (SJK).

DISCLAIMERS The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/IBC, or the U.S. Government (SJK). The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.

Methods

Animal Preparation

All experimental procedures were compliant with the guidelines established by the Institutional Animal Care and Use Committee of Carnegie Mellon University and the National Institutes of Health, and all experimental protocols were approved by the Institutional Animal Care and Use Committee of Carnegie Mellon University. To express the calcium indicator GCaMP6f selectively in excitatory neurons, homozygous Emx1cre mice (Jackson Laboratories, stock number 005628) were crossed with homozygous Ai93/heterozygous Camk2a-tTA mice (Jackson Laboratories, stock number 024108). Experimental mice were heterozygous for all three alleles. Mice were housed in groups of 2-3 per cage, in a 12-hour light/12-hour dark cycle and all imaging was performed at Zeitgeber time (ZT) 14.5±1, where ZT0 is lights on, and ZT12 is lights off. The same enrichment materials were provided in all cages including a Plexiglas hut and nesting material. Fourteen mice were imaged in this study: 4 mice were used for BCI training (2 males(M), 2 females (F)), 6 mice were used for dark exposure (4M, 2F), and 4 mice were used as controls for decoding of orientation and spatial frequency (3M, 1F). 13 out of 14 mice imaged were treated with doxycycline from conception to 2 weeks of age[30] to reduce aberrant, interictal events sometimes observed in this transgenic line of mice[43,56]. None of the mice exhibited aberrant activity in V1 and adjacent regions.

A 3-mm diameter craniotomy was made over the primary visual cortex in the left hemisphere, identified by coordinates and landmarks as described in[57]. Mice (29-39 days old) were anesthetized with isoflurane (3% induction, 1-2% maintenance). A stainless-steel bar, used to immobilize the head for recordings, was glued to the right side of the skull and secured with dental cement. The craniotomy was then covered with a double glass assembly in which the diameter of the inner glass was fitted to the craniotomy and sealed with dental cement. Mice were allowed to recover for a minimum of 3 days with ad libitum access to food and water.

Data Acquisition, Neuron Segmentation, and Neuron Tracking

Two-photon calcium imaging was performed in awake head-fixed mice mounted atop a floating spherical treadmill using a resonant scanning microscope (Neurolabware) outfitted with a 16x Nikon objective (0.80 NA). Treadmill motion was recorded using a camera (Dalsa Genie M640-1/3), and eye blinks were captured using a second camera (Dalsa Genie M1280)[43]. During BCI sessions, a separate tracking laser sensor (Keyence LV-N11MN) was used to sample the motion of the treadmill at 200 ms intervals[58].

A laser excitation wavelength of 920 nm was used (Coherent, Inc.); green emissions were filtered (Semrock 510/84-50), amplified (Edmund Optics 59-179), and detected with a PMT (Hamamatsu H1 0770B-40). The imaged field of view was 620 x 504 microns, pixel dimensions were 0.85 X 0.98 µm, and the acquisition rate was 15.5 Hz. The acquired image time series were motion-corrected by computing the horizontal and vertical translation of each frame using phase correlation[43], and individual neurons segmented using the Matlab version of Suite2p toolbox[59], as described in[43].

To identify neurons that were tracked across imaging sessions, we registered repeat imaging sessions using the mean intensity image of each session. The mean intensity image for a session was computed by averaging the intensity of each pixel in the aligned calcium image series across time for the entire imaging session (roughly 50000 frames). Then, the

mean intensity images of the two sessions were registered using an affine transform with one-plus-one evolutionary optimizer. Once the sessions were registered, the percentage of pixel overlap between the neurons from two sessions was computed. Neurons were accepted to be the same neuron across sessions, if the percentage of overlapping pixels with respect to both the segment from the first session and the segment from the second session was larger than 75%. On average, there were 160 pixels in a given neuron.

Auditory Go/No-go Pitch Association Task

Mice learned to discriminate between a 15 kHz 'Go' stimulus and a 5 kHz 'No-go' stimulus prior to BCI training. After recovery from surgery, mice were placed on water restriction. During water restriction, ad libitum access to water was removed. Each day, mice were provided with 750μL of water placed in a dish, and the dish was available for up to 30 minutes. The weights of the mice were closely monitored, and weight loss was observed during the early days of water restriction. Training started when the weight of the mice stabilized. We considered the weight of mice to be stabilized when the weight of the mice reached 80% of the weight at the start of water restriction, and the day-to-day change in weight was less than 0.1 grams for a minimum of 3 days.

Mice were head-fixed atop a spherical treadmill, with a custom-built lick port positioned near their mouth. The auditory stimuli were generated in Matlab (Mathworks, Inc.) by generating a 200 ms sinusoidal waveform of either 5kHz or 15kHz frequency with a Gaussian mask, thus creating a 200ms pulse. The sampling rate of a pulse was 44.1kHz. The generated pulses were repeated during the duration of a trial at 5Hz. The stimulus was presented through a speaker positioned 50 degrees to the left of the mouse with respect to the midline at a distance of 30 cm. Each trial started when the stimulus sounded through the speaker. Lick port was armed with reward after a 200ms sensory delay period only on Go trials. Trials lasted between 2-4 seconds, which remained fixed within a given session. The stimulus lasted for the duration of the trial, and the reward was available for retrieval while the Go stimulus was on. A single duration was used per experimental session. Licking was detected by a photodiode sensor positioned at the end of the lick port tube. One reward was delivered per trial, and each reward was about 8μL in size. Trials in which animals correctly licked during Go trials and withheld licks during No-Go trials were counted as successful trials. When a mouse licked on a No-Go trial (false alarm trials), a timeout was enforced. During the timeout, there was no auditory stimulus, and the animal had to withhold the lick for a minimum duration selected by the experimenter to exit the timeout period. Each time the animal licked during a timeout, the timeout timer was restarted. Upon exiting the timeout period, the mice immediately entered the next trial. The range of timeout used was between 3 and 8 seconds. Performance in each session was quantified as the maximum percentage of successful trials in 100 consecutive trials. Mice were trained 1 session per day and each session was 30 to 60 minutes long, during which the mice completed 200 to 500 total trials. After reaching 80% performance, mice were advanced to BCI training.

Brain Computer Interface Task

Mice were trained on a brain-computer interface task in which they could earn rewards by modulating the activity of six neurons to control an auditory pitch. We refer to these neurons

as the direct neurons. The six neurons were randomly selected from a pool of neurons that were tracked and well isolated from other neurons, as such not all neurons were visually responsive to grating stimuli. The neural activity of the direct neurons was transformed into an auditory pitch using the following method: Prior to each BCI session, we measured the spontaneous fluorescence activity of the direct neurons for 3 minutes in the dark (3000 frames). The mean and standard deviation of this spontaneous fluorescence activity was then used to normalize the real-time raw activity into z-scores. During BCI control, the real-time direct neuron activities were z-score normalized and binned into 200 ms bins. Then, the activities of the six neurons were transformed into a one-dimensional control signal that we defined as the neural drive. The neural drive at every time bin t was calculated according to the equation below.

$$\text{Neural Drive} = \sum_{i=1}^{3} Z_i - \sum_{j=1}^{3} Zj$$

In the equation above, Z is the z-scored activity of a given neuron and i and j is the index of the neuron. According to the equation, three of the neurons contributed positively to the neural drive such that an increase in the activity of these neurons increased the neural drive, while the z-scored activity from the other three neurons contributed negatively to the decoder so that an increase in the activity of these neurons decreased the neural drive. We define neurons that positively contribute to the neural drive as direct positive neurons (DPs) and neurons that negatively contribute to the neural drive as direct negative neurons (DNs). Neural drive values were mapped to discrete target pitches as follows: neural drive values < 20% of the target threshold were mapped to 5kHz, values between 20% and 40% were mapped to 6.23kHz; values between 40% and 60% were mapped to 6.23kHz; values between 60% and 80% were mapped to 9.67kHz; values between 80% and 100% of the target were mapped to 12.04kHz; and values greater than or equal to the target were mapped to 15.00kHz. Once BCI training was initiated, no more than 2-3 consecutive days occurred between BCI training sessions.

Time of success was defined as the bin in which the auditory pitch reached 15 kHz. Mice had up to 10 seconds to reach the target once the trial started. Mice then had up to 4 seconds to retrieve the water reward by licking the lick port. Only a single drop (~8µL) was delivered per success. Prior to trial start, we enforced an inter-trial-interval where the mice had to satisfy two conditions in order to start a trial. First, mice could not move. Second, the neural drive had to be lower than the target threshold. When both of these conditions had been met for 5 consecutive time bins (1 second), the trial was started. The trial start was cued to the animal by the onset of auditory feedback. The target for each mouse was determined so that the mouse would have succeeded in about 30% of the trials during a 10-minute recording of spontaneous activity in the dark, recorded before mice were ever exposed to BCI training. In all BCI sessions within a mouse, the same six neurons were used, and the same target was used across the sessions. However, the mean and standard deviation of the fluorescent activity used to z-score the activity of the direct neuron was re-measured in each session to control for variations in imaging conditions across sessions.

Visual Stimulation

Static sinusoidal grating stimuli were generated using psychophysics toolbox (http://psychtoolbox.org/) in Matlab (Mathworks, Boston, MA). The stimulus was presented on a

screen positioned 25 cm away from the right eye angled at 50° with respect to the midline of the animal. The size of the screen was 64 cm by 40 cm, thereby subtending 142° × 96° of visual angle. The spatial frequency range of the stimulus set was 0.02 cycles/° to 0.3 cycles/° at 0.02 cycles/° interval. The orientations ranged from 0° to 180° at 15° spacing interval, yielding a total of 180 different sinusoidal gratings with 12 different orientations and 15 different spatial frequencies. Each grating was presented for 250 ms consecutively in a random order without interleaved gray screen. Each stimulus was shown at least 20 times.

Quantification of Visual Responses

Reverse correlation was used to determine the onset of cortical response[60]. The delay of the analysis window relative to the onset of the stimulus was determined by computing the timing of the peak in the stimulus-averaged events. The peak in the stimulus-averaged events was observed 194 ms to 320 ms after the stimulus was presented on the screen. Therefore, for each stimulus, the corresponding event activity was computed by averaging the number of events between 194 ms and 320 ms window. We defined this period as the response window for a given stimulus.

A neuron was defined as responsive to visual stimuli when the number of events following a presentation of a visual stimulus was modulated by the stimuli presented. To test for modulation, we performed a one-way analysis of variance (ANOVA) on the observed events during the response window using stimuli as the factor for each neuron. A neuron was considered to be responsive when the result of ANOVA was significant at 99% significance level.

GCaMP6f expressed in neurons had longer decay than the presentation rate of our stimuli[61]. Therefore, we used deconvolution to remove the effects of decay in calcium fluorescence in quantifying responses of each neuron to our visual stimuli as in[30]. Briefly, amplitude of calcium transients was expressed in units of inferred events. For each segment n, inferred events s_n were estimated from fluorescence using the following model:

$$f_n = s_n * k + \beta_n p_n + b_n$$

where k is the temporal kernel and b_n is the baseline fluorescence. Neuropil, which is a contamination of the fluorescence signal f_n from out of focus cell bodies and nearby axons and dendrites, is modeled by p_n, the time course of the neuropil contamination, and, β_n the scaling coefficients. * denotes convolution. Using this model, s_n, k, β_n, and b_n were estimated by a matching pursuit algorithm with L0 constraint, in which spikes were iteratively added and refined until the threshold determined by the variance of the signal was met.

Trials containing locomotion or eye blinks were removed. Pupil location was estimated from eye-tracking videos using a circular Hough Transform algorithm; the algorithm failed to find the pupil on frames during which the mice were blinking. These frames were marked as eye blink frames and removed from further analysis. Trials with locomotion were identified as in[30]. Briefly, after applying a threshold on the luminance intensity of the treadmill motion images, phase correlation was computed between consecutive frames to estimate the translation between the frames. To define a motion threshold, the data were smoothed using a 1s sliding window. Any continuous non-zero movement periods during which the animal's instantaneous running speed exceeded 10 cm/s threshold for at least one frame were marked as running epochs.

Estimation of preferred stimulus and tuning bandwidth (*Figures 3,4*)

Orientation and spatial frequency preference were determined using a two-dimensional Gaussian model, fit to single trial responses. For neurons that were responsive to grating stimuli, a two-dimensional Gaussian model was fit using nonlinear least-squared regression such that the number of events R as a function of the orientation θ and the spatial frequency φ of the stimulus was

$$R(\theta,\varphi) = \frac{A}{2\pi\sigma_\theta\sigma_\varphi\sqrt{1-\rho^2}} e^{\left(-\frac{1}{2(1-\rho)^2}\left[\frac{(\theta-\mu_\theta)^2}{\sigma_\theta^2}+\frac{(\varphi-\mu_\varphi)^2}{\sigma_\varphi^2}-\frac{2\rho(\theta-\mu_\theta)(\varphi-\mu_\varphi)}{\sigma_\theta\sigma_\varphi}\right]\right)} + B$$

where μ_θ was the preferred orientation and μ_φ was the preferred spatial frequency of the stimulus, and the σ_θ and σ_φ described the widths of respective tuning. The covariance of responses for orientation and spatial frequency was captured by the correlation term ρ. A was a parameter accounting for the amplitude of the responses in number of events, while B was the baseline event activity of the cell. For fitting, the lower and the upper bound of allowed values for μ_φ was set by the range of the presented stimuli, which was 0.02 to 0.30 cycles/°. The lower bound for σ_θ and σ_φ was set at 1 degree and 0.001 cycles/° respectively to prevent fits with zero or negative widths. Prior to fitting, the preferred orientation was initialized by estimating the preferred orientation by averaging the response, *R* across all spatial frequencies for a given stimulus orientation, θ and calculating half the complex phase of the value

$$S = \frac{\sum R(\theta)e^{2i\theta}}{\sum R(\theta)}$$

[32,62]. The preferred spatial frequency was initialized by selecting the spatial frequency that generated the maximal significant response at the estimated preferred orientation.

For the model above, R^2 of the fit was used to find neurons with significant tuning. The chance distribution of R^2 was calculated from fitting the above model with permuted stimulus labels on individual trials 1000 times for each neuron. Neurons whose R^2 exceeded the 95th percentile of the chance R^2 distribution were accepted as tuned to grating stimuli.

The bandwidths of the Gaussian tuning were described using half-width at half-maximum (HWHM). The HWHM bandwidths for both orientation and spatial frequency were calculated as

$$\text{BW} = \sqrt{2 * \ln(2)} * \sigma$$

where σ was the width parameter of the Gaussian fit.

Reliability and amplitude quantification (*Figure 3*)

For each neuron that were responsive and tuned, we determined the reliability and amplitude of response to grating stimuli[63]. To do so, we compared the measured number of deconvolved events in response to each presentation of a stimulus to a null distribution of event activity in the absence of a stimulus for each neuron. Then, on a neuron-by-neuron basis, the null distribution was built from the spontaneous event activity collected while the screen was off. The collected spontaneous event activity was binned into same size bins as the stimulus response window (2 calcium imaging frames). This binned spontaneous event activity was then used as the null distribution. The number of samples for the null distributions was matched

across all imaging sessions as the minimum number of bins available across the sessions. The minimum number of samples was 776. For sessions with more than the minimum samples, 776 were randomly selected from available samples.

After we computed the null distribution for each neuron, we compared response from each presentation of a stimulus to the computed null distribution to identify trials in which a significant response was detected. We labelled trials with the number of events larger than 95% of the null distribution as trials in which the responses were significant. Then for each stimulus and each neuron, we calculated the reliability as the number of trials with significant responses over the total number trials after removing running and blinking trials. In addition, we computed the response amplitude to each stimulus as the average number of deconvolved events across trials with significant response.

Decoding (*Figures 3,4*)

K-nearest-neighbor classifiers were used to decode orientations from vectors of single trial population responses to grating stimuli presentation[43,63]. In our case, the k-nearest-neighbor classifier estimated the stimulus identity for a given response vector by identifying the most frequent stimulus identity of its k closest response vectors. To identify the nearest neighbors for a given response vector, r_t, we used correlation distance to other response vectors, r_n:

$$\text{Correlation Distance} = 1 - \text{corr}(r_t, r_n)$$

For each session, data was divided so that a single set of response vectors consisted of one trial of each stimulus. This resulted in the number of sets being equal to the number of trials that each stimulus was shown. During decoding, the possible neighbors for a test response vector consisted of all response vectors not belonging to the test set. This ensures an unbiased representation of possible nearest neighbors across stimuli. This process was repeated across each response vector and each set. We reported the performance of this decoding process as accuracy across all response vector tested. Only the neurons that were responsive to grating stimuli were included in decoding. To compare accuracies across sessions, we matched the number of neurons used in classification to the minimum number of neurons across the session pair. To ensure that the computed accuracies were not biased by subsampling of neurons, each session was decoded 10 times using a randomly selected set of neurons, and we computed the average accuracy for each session. To determine the optimal value of k, the number of neighbors, we performed decoding on the first session of each mouse sweeping k from 3 to 15. Note that this session is not used for other analyses, ensuring a properly cross-validated estimate of k. Three of four mice had accuracy ranges greater than 5% across different values of k. For the three mice that had accuracy ranges greater than 5% across different values of k, we ranked the values of k that yielded the highest accuracy for each mouse. We found k = 4 had the best average rank across mice. Therefore, we used k = 4 to decode orientations for sessions that were compared in the main results. The chance performance of the decoder was 8.33%.

Computation of signal correlations (*Figure 4*)

Signal correlation ρ^{sig} between a pair of neurons is defined as Pearson's correlation between the average responses to stimuli[64]. Therefore, we computed pairwise signal correlation between neuron i and neuron j as

$$\rho_{i,j}^{sig} = \mathrm{corr}(\bar{R}_i, \bar{R}_j)$$

where $\bar{R}$ is a vector of average response in number of spikes to 180 sinusoidal gratings for the respective neuron.

BCI Data Analysis

<u>Quantification of excitability at time of success (*Figure 2*)</u>

The fluorescence signal of identified neurons was normalized post-hoc (z-score). The post-hoc z-score normalization was performed in an identical manner to the real-time normalization of direct neurons, including binning the signal into 200 ms blocks (each bin represented the sum of three imaging frames), except that the fluorescence signal was extracted using Suite2p. For each successful trial, the signal amplitude at the time of success was used to build a distribution of activity for a given neuron. For each neuron that was tracked in the first session and learning point session, we tested whether the distribution of activity at the time of success was shifted in the learning point session compared to the first session using Wilcoxon rank-sum test. Neurons that exceeded 95% significance level were defined to have a significant shift in the excitability.

<u>Definition of neural trajectories (*Figure 5A*)</u>

To define neural trajectories, we examined the contribution of direct positive neurons and direct negative neurons to reaching the target. We measured the contribution to reaching the target as the ratio of the summed activity of the direct neurons to the target. Contribution for the direct positive neurons, C_p was defined as:

$$C_p = \frac{\sum DP}{T}$$

$\sum DP$ was the summed activity of direct positive neurons at the time of success and T was the target in the neural drive space for a given mouse. Similarly, contribution for the direct negative neurons, C_n was defined by the equation below:

$$C_n = \frac{-\sum DN}{T}$$

For each success trial, the neural trajectory class was determined by the amount of C_p and C_n: if C_p and C_n both exceeded 20% of the target, a given success trial was classified as both ΣDP increased & ΣDN decreased trajectory; if only $C_p$was greater than 20% but not C_n, it was categorized as ΣDP increased trajectory; and if only $C_n$was greater than 20% but not C_p, it was categorized as ΣDN decreased trajectory. A contribution of 20% would increase the auditory pitch one step.

Statistics

Error is reported as standard error of the mean (S.E.M.), unless noted. In the case data were not normally distributed, non-parametric tests were used.

References

1. Allen, W. E. *et al.* Global Representations of Goal-Directed Behavior in Distinct Cell Types of Mouse Neocortex. *Neuron* **94**, 891-907.e6 (2017).
2. Koralek, A. C., Jin, X., Long, J. D. 2nd, Costa, R. M. & Carmena, J. M. Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature* **483**, 331–335 (2012).
3. Roth, R. H. *et al.* Cortical Synaptic AMPA Receptor Plasticity during Motor Learning. *Neuron* **105**, 895-908.e5 (2020).
4. Neely, R. M., Koralek, A. C., Athalye, V. R., Costa, R. M. & Carmena, J. M. Volitional modulation of primary visual cortex activity requires the basal ganglia. *Neuron* **97**, 1356-1368.e4 (2018).
5. Wang, Y. *et al.* A Mouse Model of Visual Perceptual Learning Reveals Alterations in Neuronal Coding and Dendritic Spine Density in the Visual Cortex. *Front. Behav. Neurosci.* **10**, 42 (2016).
6. Rose, T., Jaepel, J., Hubener, M. & Bonhoeffer, T. Cell-specific restoration of stimulus preference after monocular deprivation in the visual cortex. *Science* **352**, 1319–1322 (2016).
7. Poort, J. *et al.* Learning enhances sensory and multiple non-sensory representations in primary visual cortex. *Neuron* **86**, 1478–1490 (2015).
8. Makino, H. & Komiyama, T. Learning enhances the relative impact of top-down processing in the visual cortex. *Nat. Neurosci.* **18**, 1116–1122 (2015).
9. Keller, A. J. *et al.* Stimulus relevance modulates contrast adaptation in visual cortex. *Elife* **6**, (2017).
10. Jurjut, O., Georgieva, P., Busse, L. & Katzner, S. Learning Enhances Sensory Processing in Mouse V1 before Improving Behavior. *J. Neurosci.* **37**, 6460–6474 (2017).
11. Goltstein, P. M., Coffey, E. B. J., Roelfsema, P. R. & Pennartz, C. M. A. In vivo two-photon Ca2+ imaging reveals selective reward effects on stimulus-specific assemblies in mouse visual cortex. *J. Neurosci.* **33**, 11540–11555 (2013).
12. Henschke, J. U. *et al.* Reward Association Enhances Stimulus-Specific Representations in Primary Visual Cortex. *Curr. Biol.* **30**, 1866-1880.e5 (2020).
13. Gilbert, C. D. & Li, W. Top-down influences on visual processing. *Nat. Rev. Neurosci.* **14**, 350–363 (2013).
14. Clancy, K. B. & Mrsic-Flogel, T. D. The sensory representation of causally controlled objects. *Neuron* (2020) doi:10.1016/j.neuron.2020.12.001.
15. Clancy, K. B., Koralek, A. C., Costa, R. M., Feldman, D. E. & Carmena, J. M. Volitional modulation of optically recorded calcium signals during neuroprosthetic learning. *Nat. Neurosci.* **17**, 807–809 (2014).
16. Hira, R. *et al.* Reward-timing-dependent bidirectional modulation of cortical microcircuits during optical single-neuron operant conditioning. *Nat. Commun.* **5**, 5551 (2014).
17. Prsa, M., Galinanes, G. L. & Huber, D. Rapid integration of artificial sensory feedback during operant conditioning of motor cortex neurons. *Neuron* **93**, 929-939.e6 (2017).
18. Mitani, A., Dong, M. & Komiyama, T. Brain-computer interface with inhibitory neurons reveals subtype-specific strategies. *Curr. Biol.* **28**, 77-83.e4 (2018).
19. Oby, E. R. *et al.* New neural activity patterns emerge with long-term learning. *Proc. Natl. Acad. Sci. U. S. A.* **116**, 15210–15215 (2019).
20. Iurilli, G. *et al.* Sound-driven synaptic inhibition in primary visual cortex. *Neuron* **73**, 814–828 (2012).
21. Meijer, G. T., Montijn, J. S., Pennartz, C. M. A. & Lansink, C. S. Audiovisual Modulation in Mouse Primary Visual Cortex Depends on Cross-Modal Stimulus Configuration and

Congruency. *J. Neurosci.* **37**, 8783–8796 (2017).
22. Deneux, T. *et al.* Context-dependent signaling of coincident auditory and visual events in primary visual cortex. *Elife* **8**, (2019).
23. Knöpfel, T. *et al.* Audio-visual experience strengthens multisensory assemblies in adult mouse visual cortex. *Nat. Commun.* **10**, 5684 (2019).
24. McClure, J. P. J. & Polack, P.-O. Pure tones modulate the representation of orientation and direction in the primary visual cortex. *J. Neurophysiol.* **121**, 2202–2214 (2019).
25. Ibrahim, L. A. *et al.* Cross-Modality Sharpening of Visual Cortical Processing through Layer-1-Mediated Inhibition and Disinhibition. *Neuron* **89**, 1031–1045 (2016).
26. Ko, H. *et al.* The emergence of functional microcircuits in visual cortex. *Nature* **496**, 96–100 (2013).
27. Ko, H., Mrsic-Flogel, T. D. & Hofer, S. B. Emergence of feature-specific connectivity in cortical microcircuits in the absence of visual experience. *J. Neurosci.* **34**, 9812–9816 (2014).
28. Cossell, L. *et al.* Functional organization of excitatory synaptic strength in primary visual cortex. *Nature* **518**, 399–403 (2015).
29. Carrillo-Reid, L., Yang, W., Bando, Y., Peterka, D. S. & Yuste, R. Imprinting and recalling cortical ensembles. *Science (80-.).* **353**, 691–694 (2016).
30. Jeon, B. B., Swain, A. D., Good, J. T., Chase, S. M. & Kuhlman, S. J. Feature selectivity is stable in primary visual cortex across a range of spatial frequencies. *Sci. Rep.* **8**, 15288 (2018).
31. Montijn, J. S., Meijer, G. T., Lansink, C. S. & Pennartz, C. M. A. Population-Level Neural Codes Are Robust to Single-Neuron Variability from a Multidimensional Coding Perspective. *Cell Rep.* **16**, 2486–2498 (2016).
32. Niell, C. M. & Stryker, M. P. Highly selective receptive fields in mouse visual cortex. *J. Neurosci.* **28**, 7520–7536 (2008).
33. Ajemian, R., D'Ausilio, A., Moorman, H. & Bizzi, E. A theory for how sensorimotor skills are learned and retained in noisy and nonstationary neural circuits. *Proc. Natl. Acad. Sci. U. S. A.* **110**, E5078-87 (2013).
34. Wolpaw, J. R. The negotiated equilibrium model of spinal cord function. *J. Physiol.* **596**, 3469–3491 (2018).
35. Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Cortical activity in the null space: permitting preparation without movement. *Nat. Neurosci.* **17**, 440–448 (2014).
36. Kirkpatrick, J. *et al.* Overcoming catastrophic forgetting in neural networks. *Proc. Natl. Acad. Sci. U. S. A.* **114**, 3521–3526 (2017).
37. French. Catastrophic forgetting in connectionist networks. *Trends Cogn. Sci.* **3**, 128–135 (1999).
38. Clopath, C., Bonhoeffer, T., Hubener, M. & Rose, T. Variance and invariance of neuronal long-term representations. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* **372**, (2017).
39. Sur, M. & Leamey, C. A. Development and plasticity of cortical areas and networks. *Nat. Rev. Neurosci.* **2**, 251–262 (2001).
40. Priebe, N. J. & McGee, A. W. Mouse vision as a gateway for understanding how experience shapes neural circuits. *Front. Neural Circuits* **8**, 123 (2014).
41. Hooks, B. M. & Chen, C. Circuitry Underlying Experience-Dependent Plasticity in the Mouse Visual System. *Neuron* **106**, 21–36 (2020).
42. Espinosa, J. S. & Stryker, M. P. Development and plasticity of the primary visual cortex. *Neuron* **75**, 230–249 (2012).
43. Kowalewski, N. N. *et al.* Development of Natural Scene Representation in Primary Visual Cortex Requires Early Postnatal Experience. *Curr. Biol.* (2020) doi:10.1016/j.cub.2020.10.046.
44. Kreile, A. K., Bonhoeffer, T. & Hubener, M. Altered visual experience induces instructive

changes of orientation preference in mouse visual cortex. *J. Neurosci.* **31**, 13911–13920 (2011).

45. Hubel, D. H. & Wiesel, T. N. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* **160**, 106–154 (1962).
46. Gordon, J. A. & Stryker, M. P. Experience-dependent plasticity of binocular responses in the primary visual cortex of the mouse. *J. Neurosci.* **16**, 3274–3286 (1996).
47. Smith, S. L. & Trachtenberg, J. T. Experience-dependent binocular competition in the visual cortex begins at eye opening. *Nat. Neurosci.* **10**, 370–375 (2007).
48. Pecka, M., Han, Y., Sader, E. & Mrsic-Flogel, T. D. Experience-dependent specialization of receptive field surround for selective coding of natural scenes. *Neuron* **84**, 457–469 (2014).
49. Woolley, S. C., Rajan, R., Joshua, M. & Doupe, A. J. Emergence of Context-Dependent Variability across a Basal Ganglia Network. *Neuron* **82**, 208–223 (2014).
50. Andalman, A. S. & Fee, M. S. A basal ganglia-forebrain circuit in the songbird biases motor output to avoid vocal errors. *Proc. Natl. Acad. Sci. U. S. A.* **106**, 12518–12523 (2009).
51. Warren, T. L., Tumer, E. C., Charlesworth, J. D. & Brainard, M. S. Mechanisms and time course of vocal learning and consolidation in the adult songbird. *J. Neurophysiol.* **106**, 1806–1821 (2011).
52. Tumer, E. C. & Brainard, M. S. Performance variability enables adaptive plasticity of 'crystallized' adult birdsong. *Nature* **450**, 1240–1244 (2007).
53. Duffy, A., Abe, E., Perkel, D. J. & Fairhall, A. L. Variation in sequence dynamics improves maintenance of stereotyped behavior in an example from bird song. *Proc. Natl. Acad. Sci. U. S. A.* **116**, 9592–9597 (2019).
54. Boerlin, M., Machens, C. K. & Denève, S. Predictive coding of dynamical variables in balanced spiking networks. *PLoS Comput. Biol.* **9**, e1003258 (2013).
55. Laje, R. & Buonomano, D. V. Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nat. Neurosci.* **16**, 925–933 (2013).
56. Steinmetz, N. A. *et al.* Aberrant cortical activity in multiple GCaMP6-expressing transgenic mouse lines. *eNeuro* **4**, (2017).
57. Feese, B. D., Pafundo, D. E., Schmehl, M. N. & Kuhlman, S. J. Binocular deprivation induces both age-dependent and age-independent forms of plasticity in parvalbumin inhibitory neuron visual response properties. *J. Neurophysiol.* **119**, 738–751 (2018).
58. Swain, A. D. *et al.* Visual acuity performance level is independent of locomotion. *bioRxiv* 750844 (2019) doi:10.1101/750844.
59. Pachitariu, M. *et al.* Suite2p: beyond 10,000 neurons with standard two-photon microscopy. *bioRxiv* (2017) doi:10.1101/061507.
60. Ringach, D. L., Sapiro, G. & Shapley, R. A subspace reverse-correlation technique for the study of visual neurons. *Vision Res.* **37**, 2455–2464 (1997).
61. Chen, T.-W. *et al.* Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* **499**, 295–300 (2013).
62. Kuhlman, S. J., Tring, E. & Trachtenberg, J. T. Fast-spiking interneurons have an initial orientation bias that is lost with vision. *Nat. Neurosci.* **14**, 1121–1123 (2011).
63. de Vries, S. E. J. *et al.* A large-scale standardized physiological survey reveals functional organization of the mouse visual cortex. *Nat. Neurosci.* **23**, 138–151 (2020).
64. Averbeck, B. B., Latham, P. E. & Pouget, A. Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* **7**, 358–366 (2006).

Figure 1

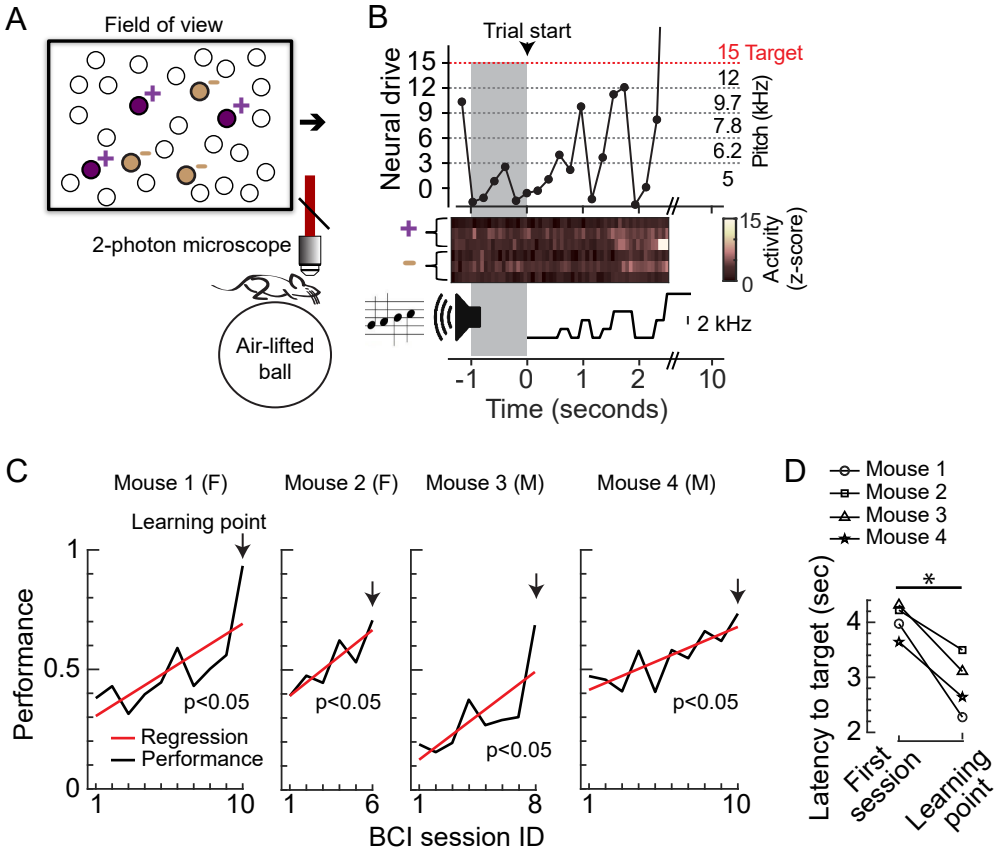

Figure 1. Mice gained volitional control of the BCI within 10 days.

(A) Schematic of BCI set-up in awake mice. Six neurons were randomly selected to drive a one-dimensional auditory cursor. In all cases, three DP neurons (purple), and three DN neurons (brown) were selected.

(B) Neural drive was mapped to the 6 auditory feedback pitch frequencies. Example trial in a trained mouse shows that the calcium signal of individual direct neurons was acquired at 15.5 Hz and accumulated for 3 frames prior to updating the pitch frequency. Pre-trial start epoch (1-second minimum, gray shading) ends when no running is detected and neural drive is less than the target threshold. The time of success is the time at which the target is reached, 2.6 seconds in this case.

(C) Fraction of trials in which the target was reached is plotted as a function of session number. The BCI session ID in which linear regression was significant is indicated as the 'learning point' (arrow; mouse 1, $p = 0.013$; mouse 2, $p = 0.023$; mouse 3, $p = 0.005$; mouse4, $p = 0.005$). M, male, F, female.

(D) The median latency to reach the target on successful trials was reduced after training, paired t-test $p= 0.011$.

Figure 2

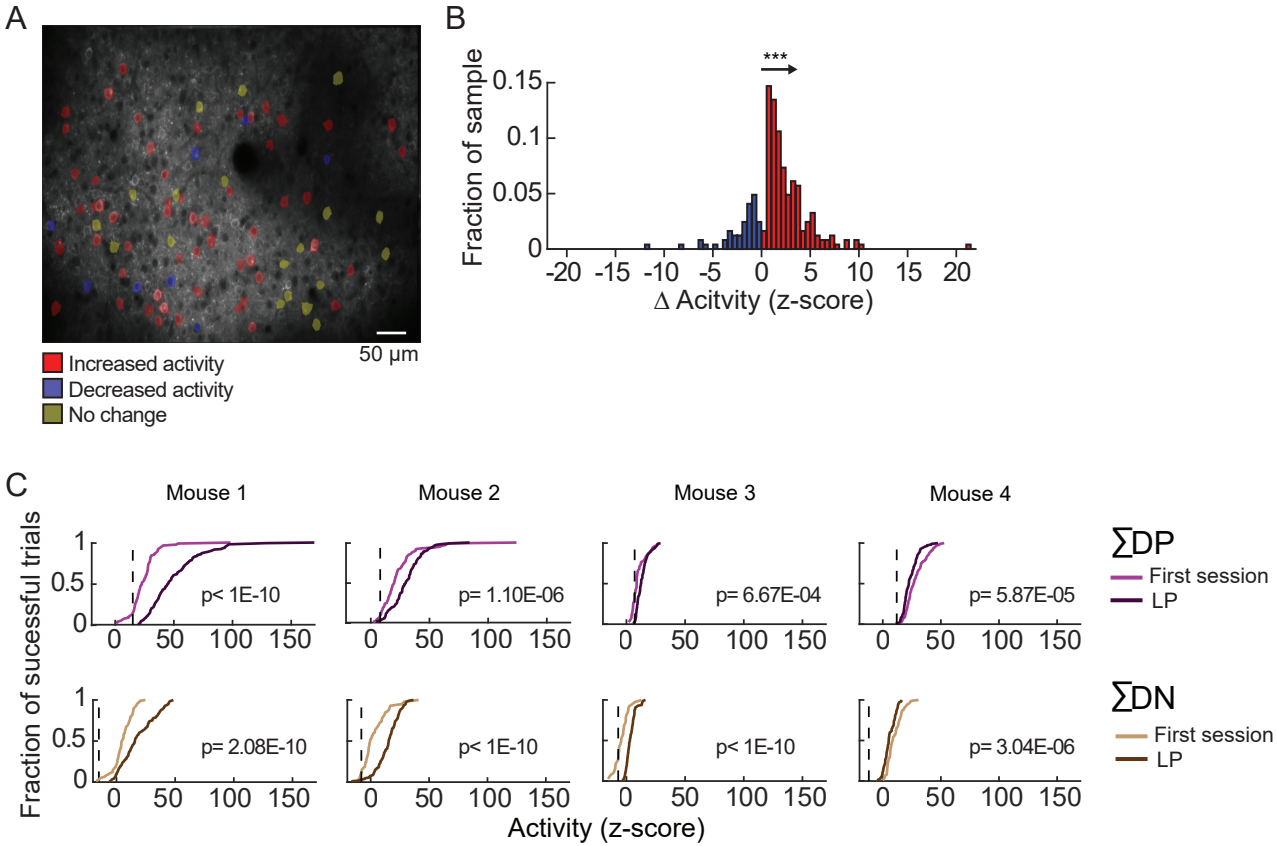

Figure 2. Volitional control of BCI is associated with a change in excitability of V1 neurons.

(A) Imaging field of view of an example animal. Indirect neurons that were tracked in Session 1 and the LP session are indicated by overlays. Changes in activity amplitude at the time of success are indicated.

(B) The distribution of the difference in activity amplitude for the neurons that exhibited a significant change between Session 1 and the LP session was skewed in the positive direction, KS-1 test, p< 1E-10.

(C) The summed (∑) activity of DP neurons was significantly higher after learning in 3 of 4 animals. In one mouse, the summed activity of DN neurons was lower after learning. KS test, p values as indicated.

Figure 3

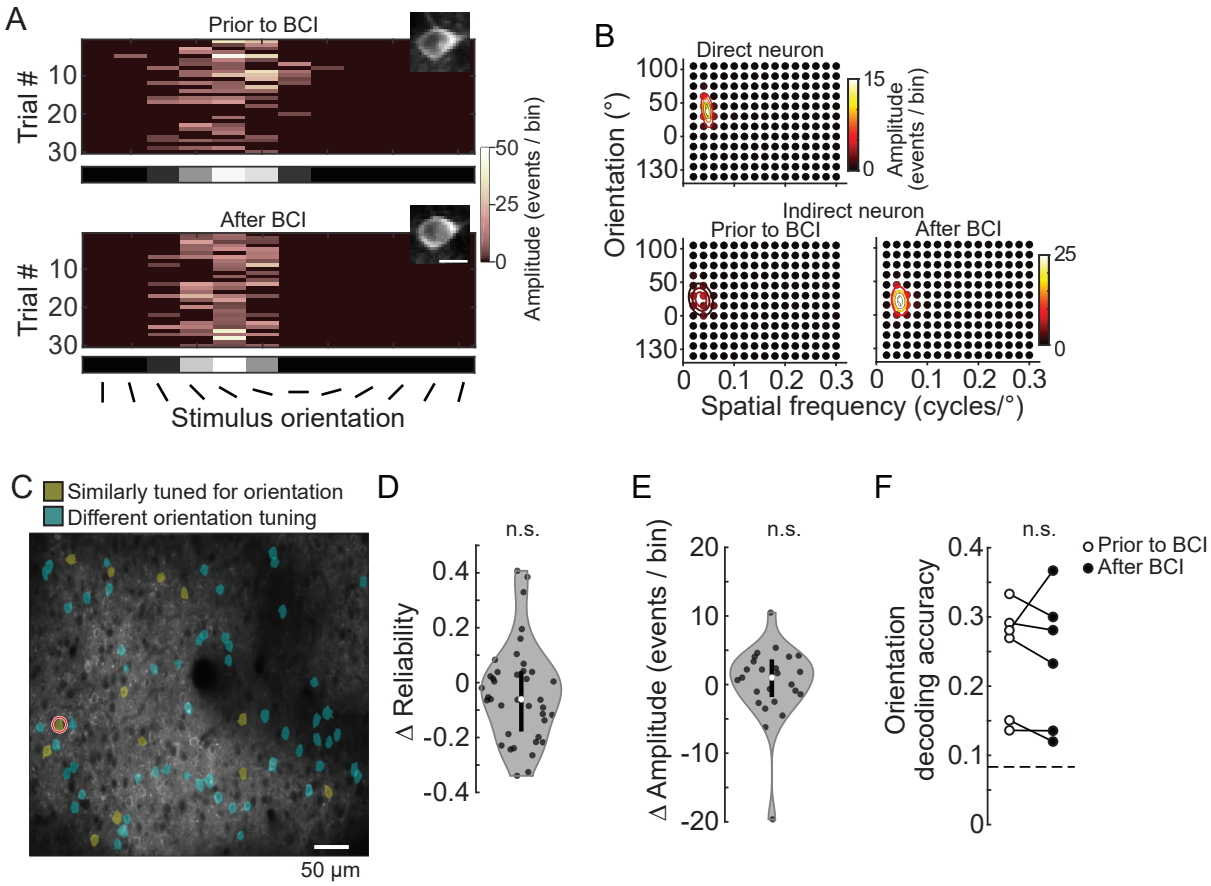

Figure 3. BCI training did not enhance stimulus responses.

(A) Example neuron responses to 12 orientations at the preferred spatial frequency, before and after BCI training. The mean across trials is shown in gray scale at the bottom (scale limits: 9, 0 events).

(B) Example 2-dimentional tuning profile of one DP neuron, and one indirect neuron before and after BCI training. Labeled contours of the Gaussian fit are 75%, 50%, and 25% of the peak response.

(C) Imaging field of view of an example animal prior to BCI training. Neurons tuned to grating stimuli both before and after BCI training are indicated by overlays. 12 indirect neurons were similarly tuned (yellow) to the direct neuron (circled in red) in this example. 57 indirect neurons were not similarly tuned (blue). Neurons were considered to be similarly tuned if their preferred orientation was within 30° and their spatial frequency within 0.03 cycles/° of a direct neuron.

(D) Difference in reliability before and after BCI training. In total across the 4 mice, 39 indirect neurons were similarly tuned to a direct neuron. The kernel density estimate(gray), median (white circle), and 25th to 75th percentile of the distribution (black bar) are indicated. The reliability of these neurons at the preferred stimulus was the same before and after BCI training, the median Δ reliability = -0.06; Wilcoxon signed-rank test on raw values, $p = 0.093$.

(E) Difference in response amplitude at the preferred stimulus before and after training. Labels as in 'D'. The response amplitude (failure trials removed) was the same before and after BCI training, the median Δ amplitude = 1.02 events/bin; Wilcoxon signed-rank test on raw values, $p = 0.382$.

(F) Decoding accuracy of orientation at the preferred spatial frequency of each tuned DP neurons, before and after BCI training. In total across the 4 mice, 6 DP neurons were tuned. The decoding accuracy was not altered by BCI training, the mean Δ orientation decoding accuracy = -0.004; paired t-test, $p = 0.936$. The number of neurons used to classify orientation was matched across condition but varied by animal: 69,33,57,47.

Figure 4

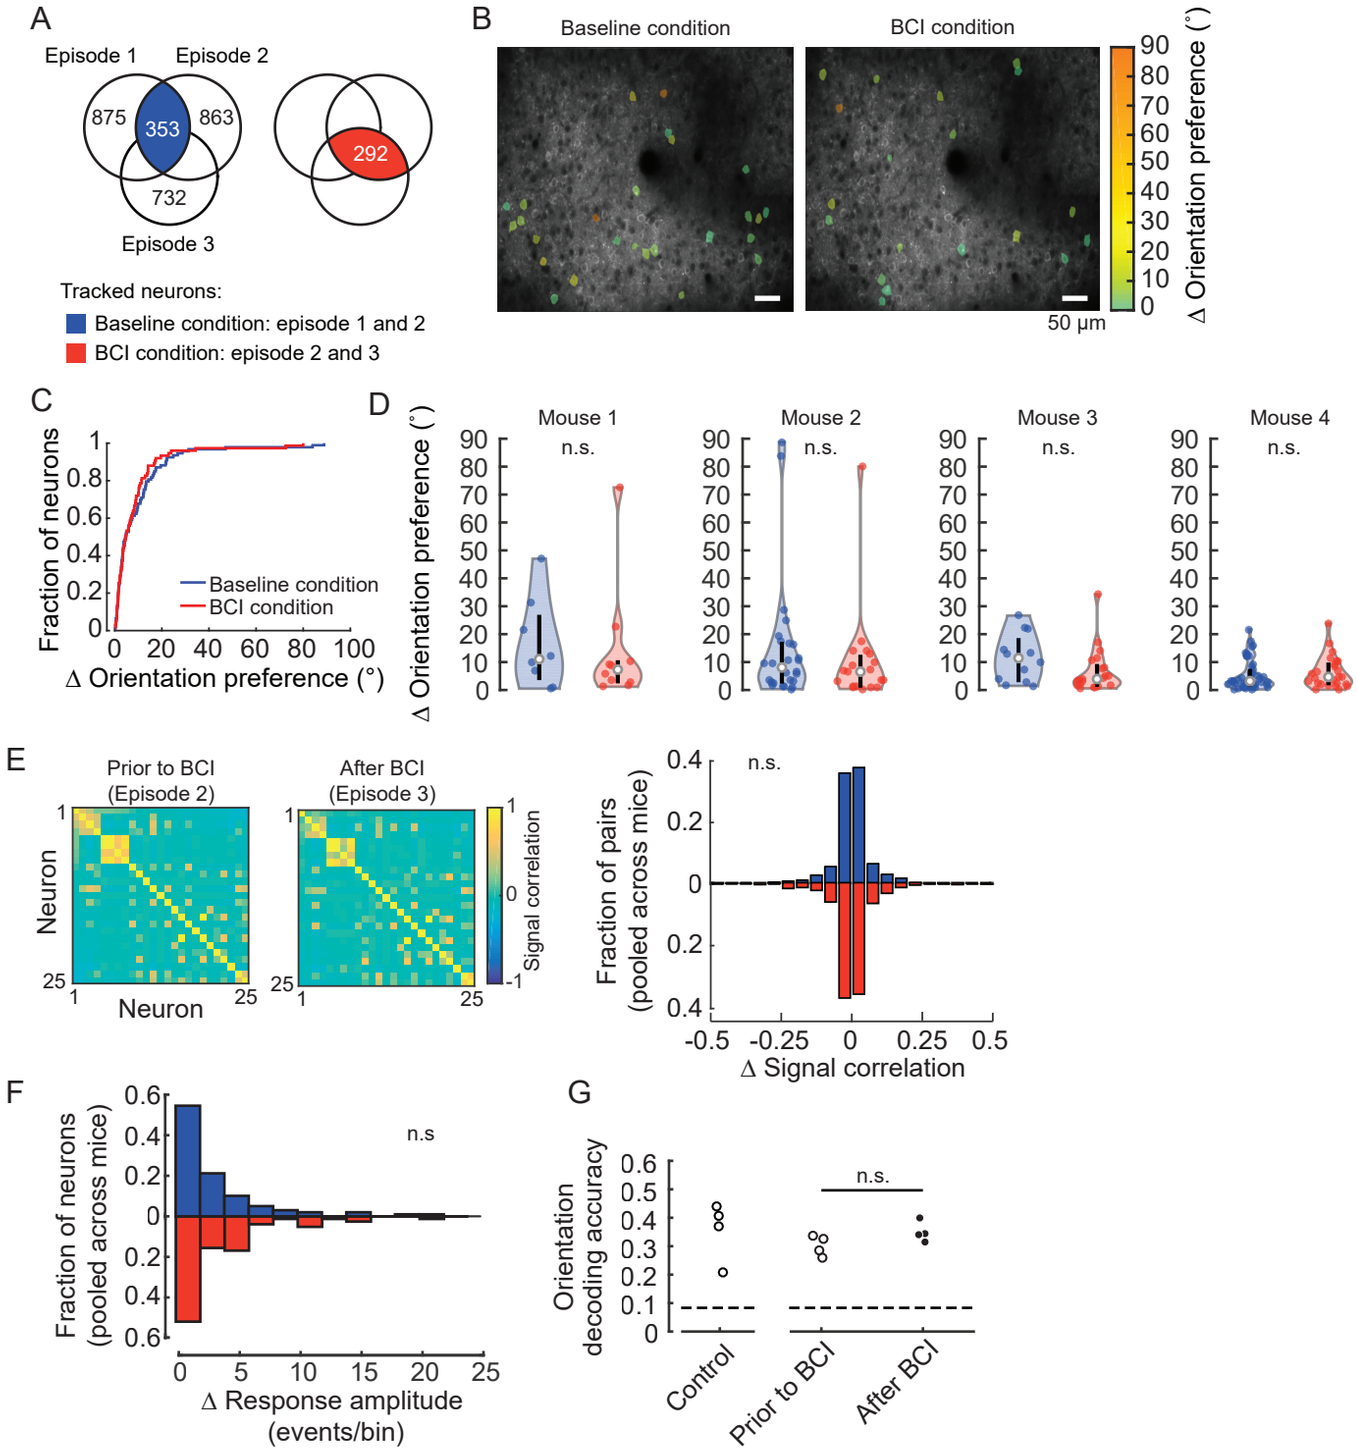

Figure 4. Acquisition of BCI control did not disrupt stimulus representation in V1.

(A) Venn diagrams depicting that across all 4 mice, a total of 353 neurons were imaged in both episode 1 and 2 (Baseline condition), and 292 neurons were imaged in both episode 2 and 3 (BCI condition). The number of imaged neurons in each episode is indicated. Episodes 1 and 2 were imaged prior to BCI training.

(B) Imaging field of view of an example animal prior to BCI training. The change in orientation preference of tracked and tuned neurons is indicated by overlays.

(C) The distribution of differences in orientation preference in the baseline condition was similar to that of the BCI condition, KS test, p = 0.680.

(D) Stability of orientation preference in the baseline condition (blue) and BCI condition (red). Labels as in Fig. 3D. Distributions were overlapping in all 4 mice; Wilcoxon rank-sum test p values: mouse 1, 8 neurons baseline, 10 neurons BCI, p = 0.633; mouse 2, 26 neurons baseline, 19 neurons BCI, p = 0.308; mouse 3, 12 neurons baseline, 21 neurons BCI, p = 0.089; mouse 4, 47 neurons baseline, 25 neurons BCI, p = 0.338.

(E) Left, Example signal correlation matrices, for the BCI condition. Right, stability of signal correlation was similar in baseline and BCI conditions, Wilcoxon rank-sum test, 1500 pairs baseline, 726 pairs BCI, p = 0.538.

(F) The difference in response amplitude in the baseline and BCI conditions was similar, Wilcoxon rank-sum test, 99 neurons baseline, 77 neurons BCI, p = 0.247).

(G) Decoding accuracy of orientation at a spatial frequency of 0.04 cycles/° was similar before (episode 2) and after (episode 3) BCI training, Wilcoxon signed-rank test, p = 0.375. The number of neurons used to classify orientation was matched across condition but varied by animal: 69,33,57,47. A second control group is plotted on the left, 69 neurons were used. Chance is 8.33% (dashed line).

Figure 5

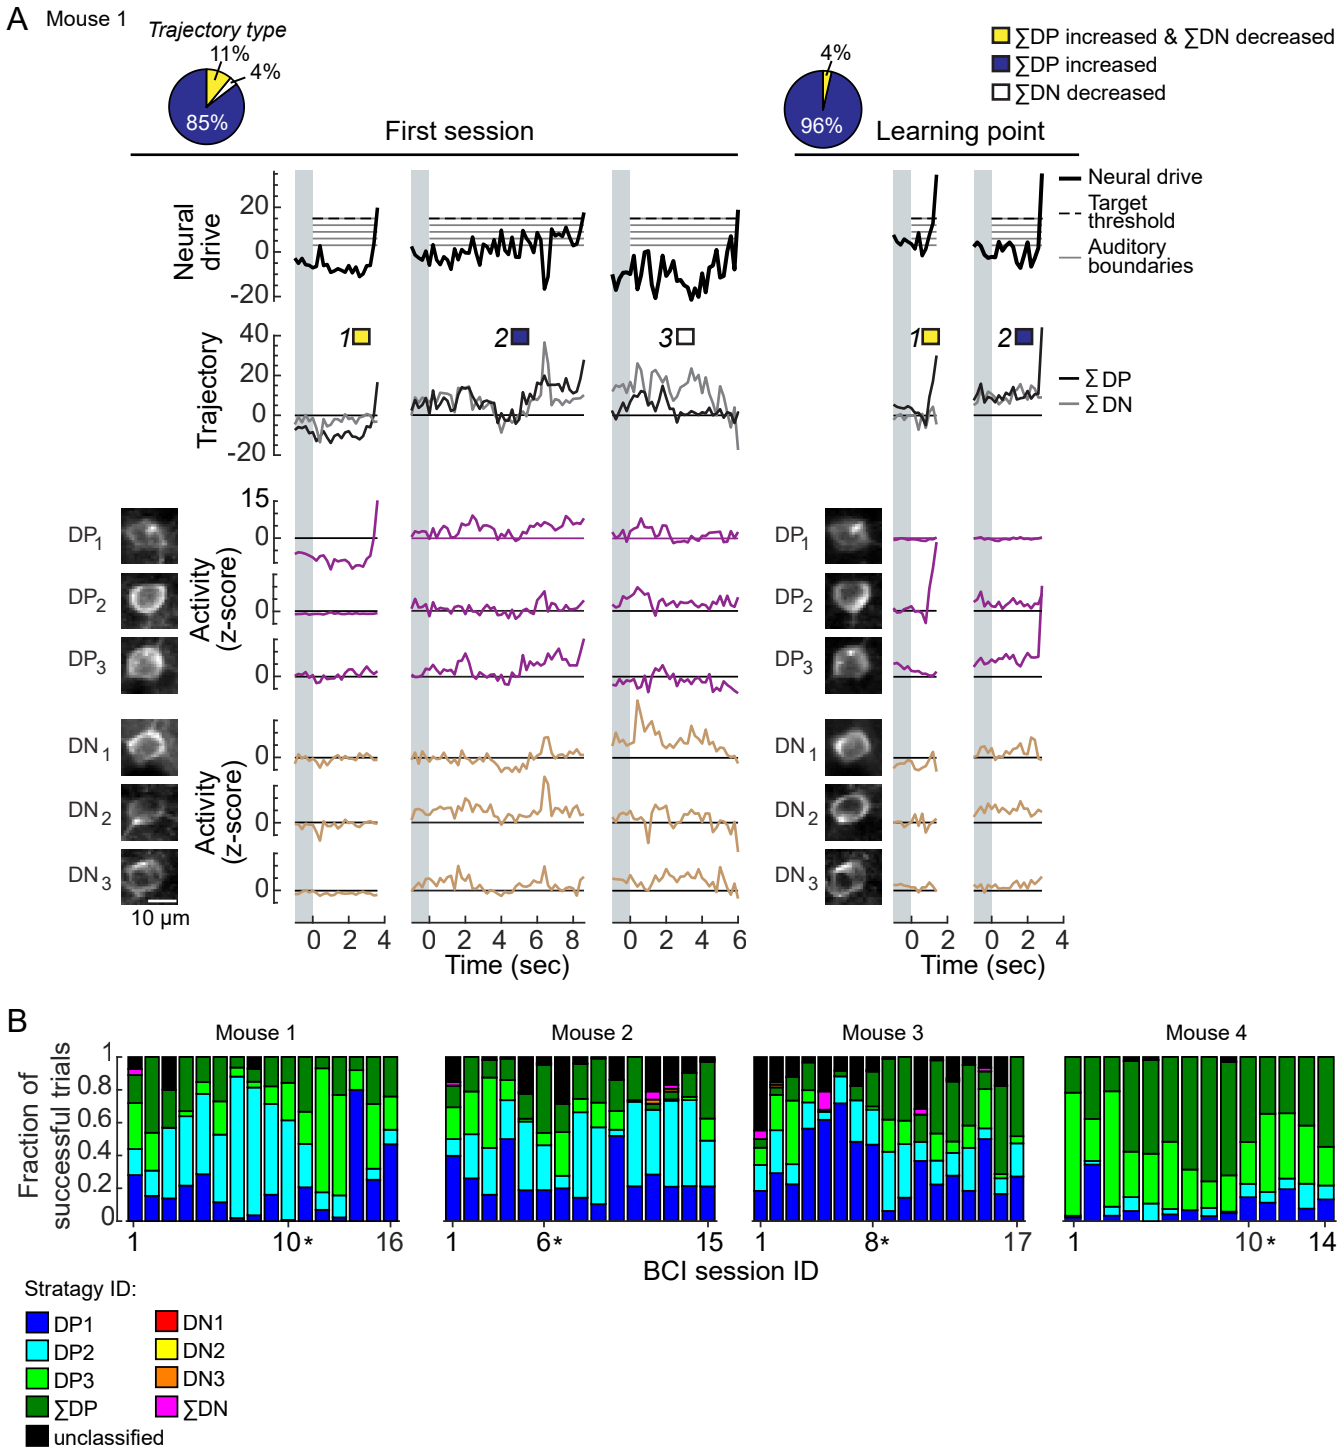

Figure 5. Flexible neural strategies were maintained after skill acquisition.

(A) Three example trials from session 1 (first 3 columns), and two example trials from the learning point session (last two columns). Frequency of trajectory types used in session 1 and the learning point session are summarized in pie charts. The same direct neurons were used throughout the experiment, DP 1-3 (purple) and DN 1-3 (brown). The six auditory pitches are indicted (gray, top row).

(B) Multiple strategies were used within a single BCI session, and across sessions. The learning point session is indicated (*).

Supplement Figure 1

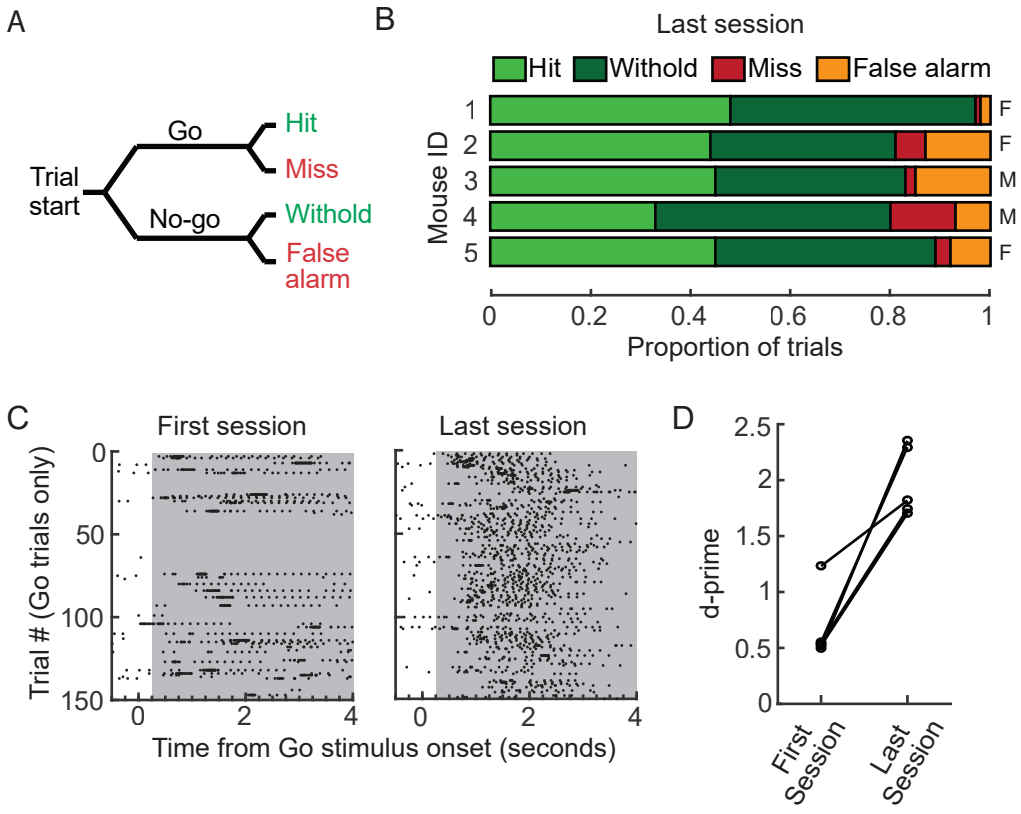

Supplement Figure 1. Mice learned to associate a 15kHz pitch with reward prior to BCI training

(A) Schematic depicting of four possible trial outcomes of the Go/No-go auditory discrimination shaping task. Correct responses are shown in green and incorrect responses are shown in red. The 'Go' stimulus was a 15kHz pitch and 'No-go' stimulus was a 5kHz pitch. Mice performed 200 to 500 trials in one session per day; one mouse was removed because it consistently failed to perform at least 50 trials per session. It is our experience that a minimum of 75 trials per session is required for mice to learn associations in Go/No-tasks.

(B) The proportion of trial outcomes for each of the five mice on the last session. Session performance is defined as the highest proportion of correct behavioral responses (hits and withholds, shown as green and dark green) calculated over a sliding window of 100 consecutive trials. Five mice trained on this task reached the 80% performance criteria and moved on to BCI training: mouse 1, 97%; mouse 2, 0.81%; mouse 3, 83%; mouse 4, 80%; mouse 5, 89%. Sex is indicated on the right: M, male, F female.

(C) Raster of lick times aligned to the onset of the 'Go' stimulus for an example mouse on the first and last session. Reward became available after a 250 ms delay from the onset of the 'Go' stimulus and remained available for the duration of stimulus presentation (gray). Mice were free to lick more than once on 'Go' trials but received only one water drop per trial.

(D) Discrimination accuracy of five mice in the first and last session reported as d-prime. All five mice had d-prime of at least 1.68 on the last session.

Supplement Figure 2

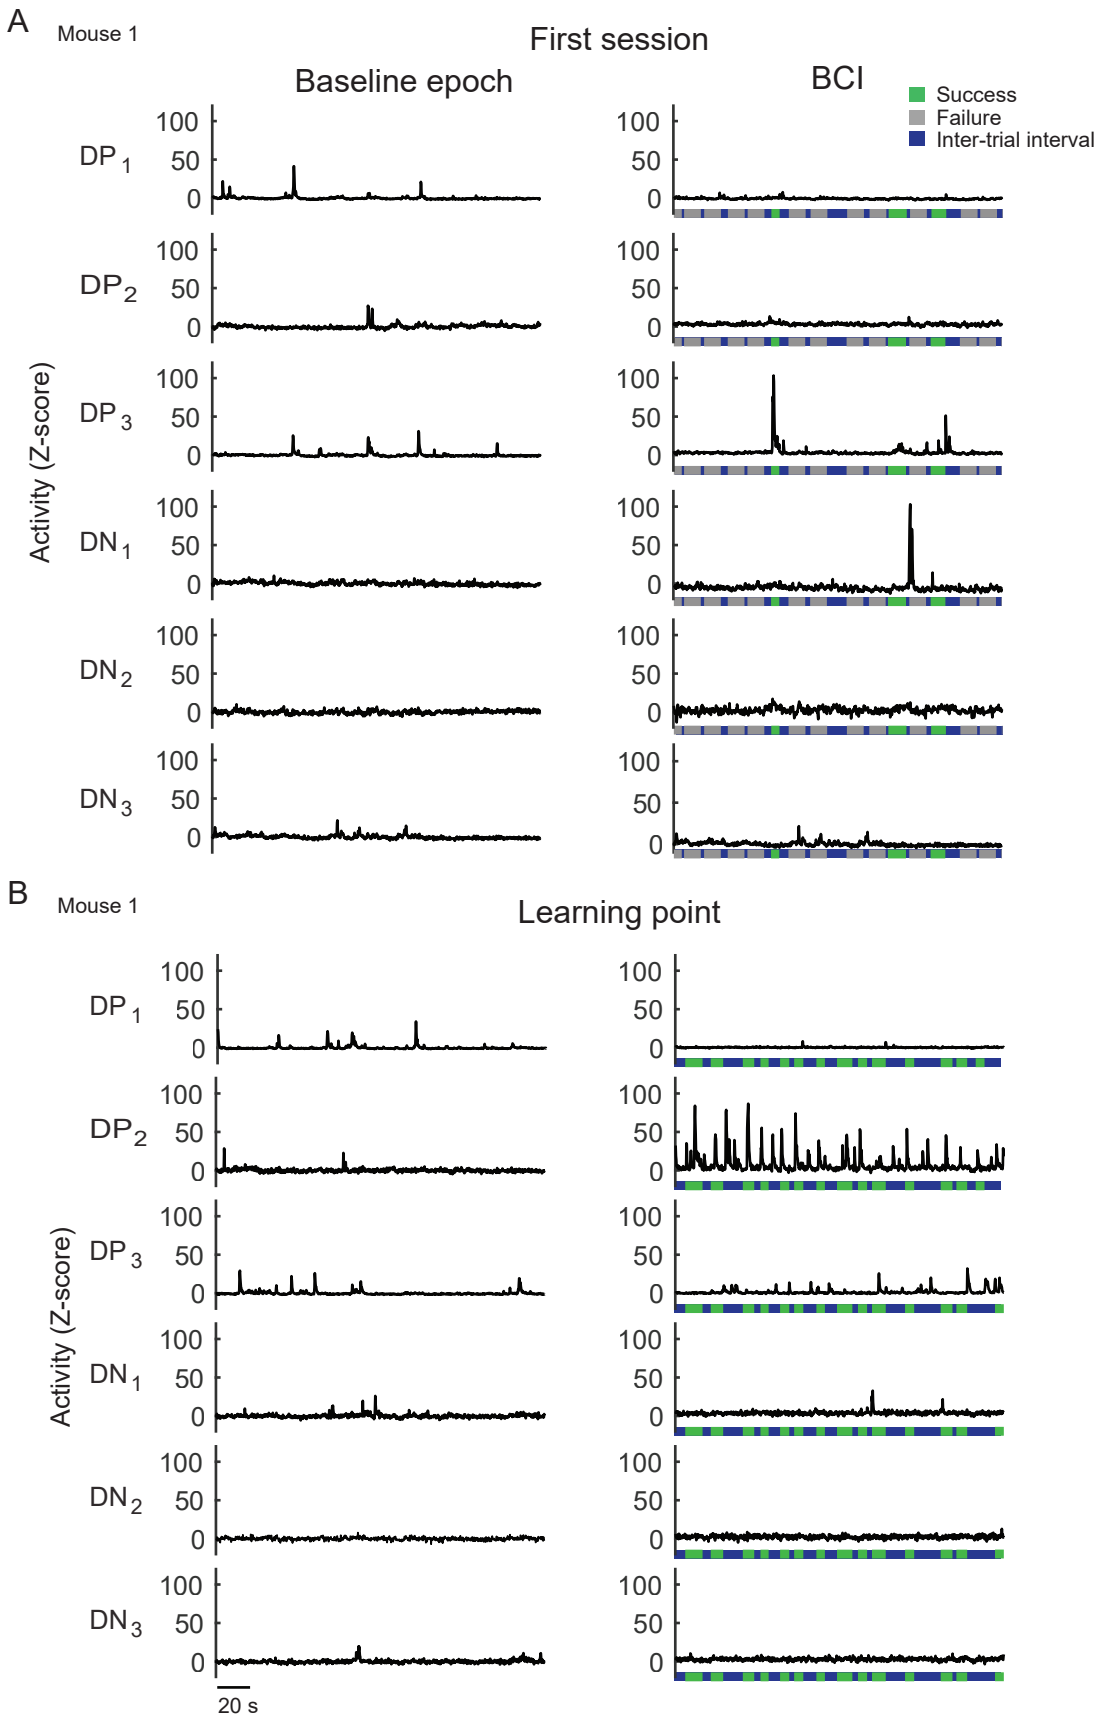

Supplement Figure 2. Examples of direct neuron activity
Example activity from one mouse, depicting 200 seconds of continuous direct neuron activity during the baseline epoch imaged immediately before BCI and during BCI training.

Supplement Figure 3

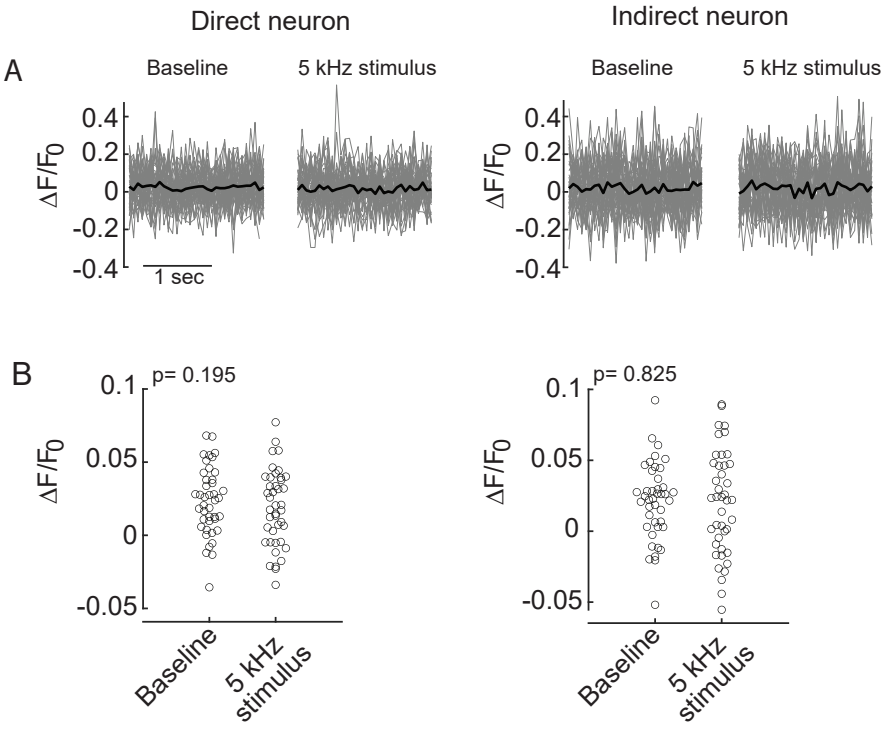

Supplement Figure 3. V1 neurons were not modulated by auditory pitches experienced during BCI training.

After BCI learning (8±6 days after the learning point session), 3 of 4 mice were presented with the 6 auditory pitches used during BCI training. Auditory pitches were presented in a random order for 2 seconds per trial, and each trial was interleaved with 2.5 seconds of silence. Each pitch was presented at least 30 trials. To test for modulation, for each pitch, the mean activity of a given neuron over the duration of a trial was compared to the mean activity during the 2-second period preceding the stimulus onset (paired t-test). We defined this 2-second duration preceding the stimulus onset as the baseline.

(A) Example activity during baseline and sound stimulus, one direct neuron (42 trials, gray) and one indirect neuron (42 trials, gray). The mean across all trials is shown in black.

(B) No difference in activity between baseline and stimulus was detected (paired t-test) for the example neurons shown in 'A'. Across the three mice, less than 3% of all neurons imaged had significant modulation (14 out of 495 neurons imaged), and none of the direct neurons were significant. A closer inspection of the few neurons that tested significant revealed that although significant, the magnitude of the difference in $\Delta F/F_0$ between baseline and sound stimulus was less than 0.03 (range: -0.02 to 0.03). Given the scarcity of the number of neurons significant and the minimal change in the amplitude for the significant neurons, we conclude that V1 neurons were not modulated by auditory pitches after BCI training.

Supplement Figure 4

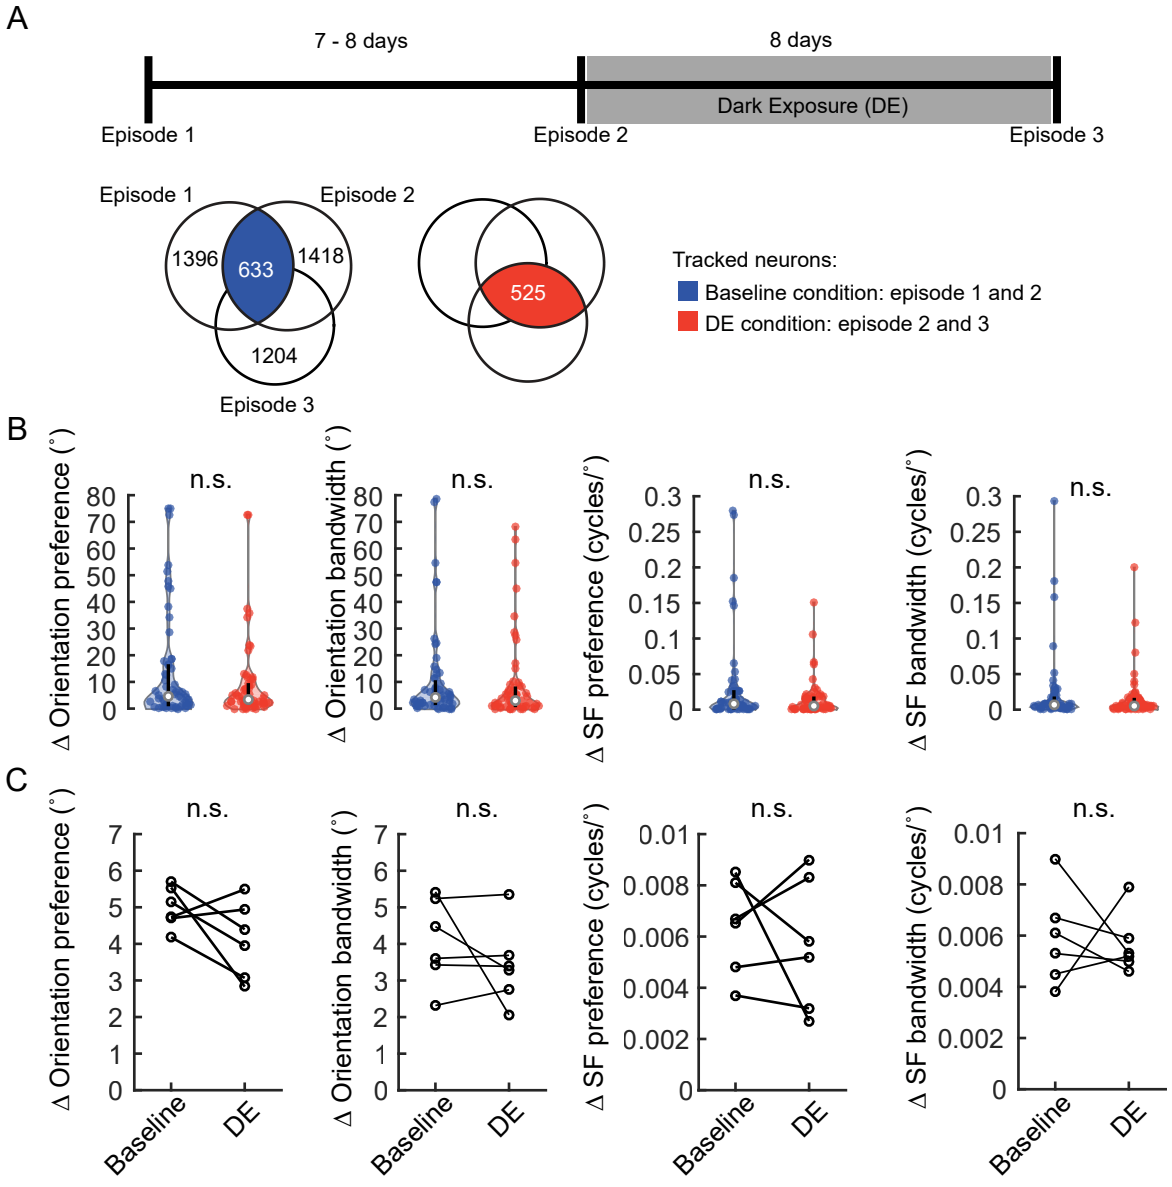

Supplement Figure 4. 8-days of dark exposure did not induce changes in the tuning parameters of V1 neurons

(A) Diagram of Dark exposure (DE) experimental design (top), and Venn diagram, labeled as in Fig. 4A (bottom). Six adult mice (x males, x females), postnatal day 45 - 86 on episode 1, were imaged in a pair of sessions 7 – 8 days apart with 12/12h light cycle. Upon completion of the episode 2 imaging session, mice were immediately moved to a dark cabinet and were kept in the cabinet for 8 days. To compare the stability of the tuning parameters before and after DE, we identified tracked neurons between episode 1 and episode 2 (baseline condition), and tracked neurons between episode 2 and episode 3 (DE condition).

(B) Stability of tuning features in an example mouse were similar in the baseline and DE conditions. Wilcoxon rank-sum test (n=60 neurons baseline condition, n=61 neurons DE condition), orientation preference: $p = 0.476$, orientation bandwidth: $p = 0.476$, spatial frequency preference: $p = 0.357$, spatial frequency bandwidth: $p = 0.532$. Labels as in Figs. 3D, 4D.

(C) Stability of tuning features in all six mice was similar in baseline and DE conditions, paired t-test, orientation preference: $p = 0.139$, orientation bandwidth: $p = 0.311$, spatial frequency preference: $p = 0.602$, spatial frequency bandwidth: $p = 0.823$. Each point is the mean difference across tuned neurons for a given mouse.

Supplement Figure 5

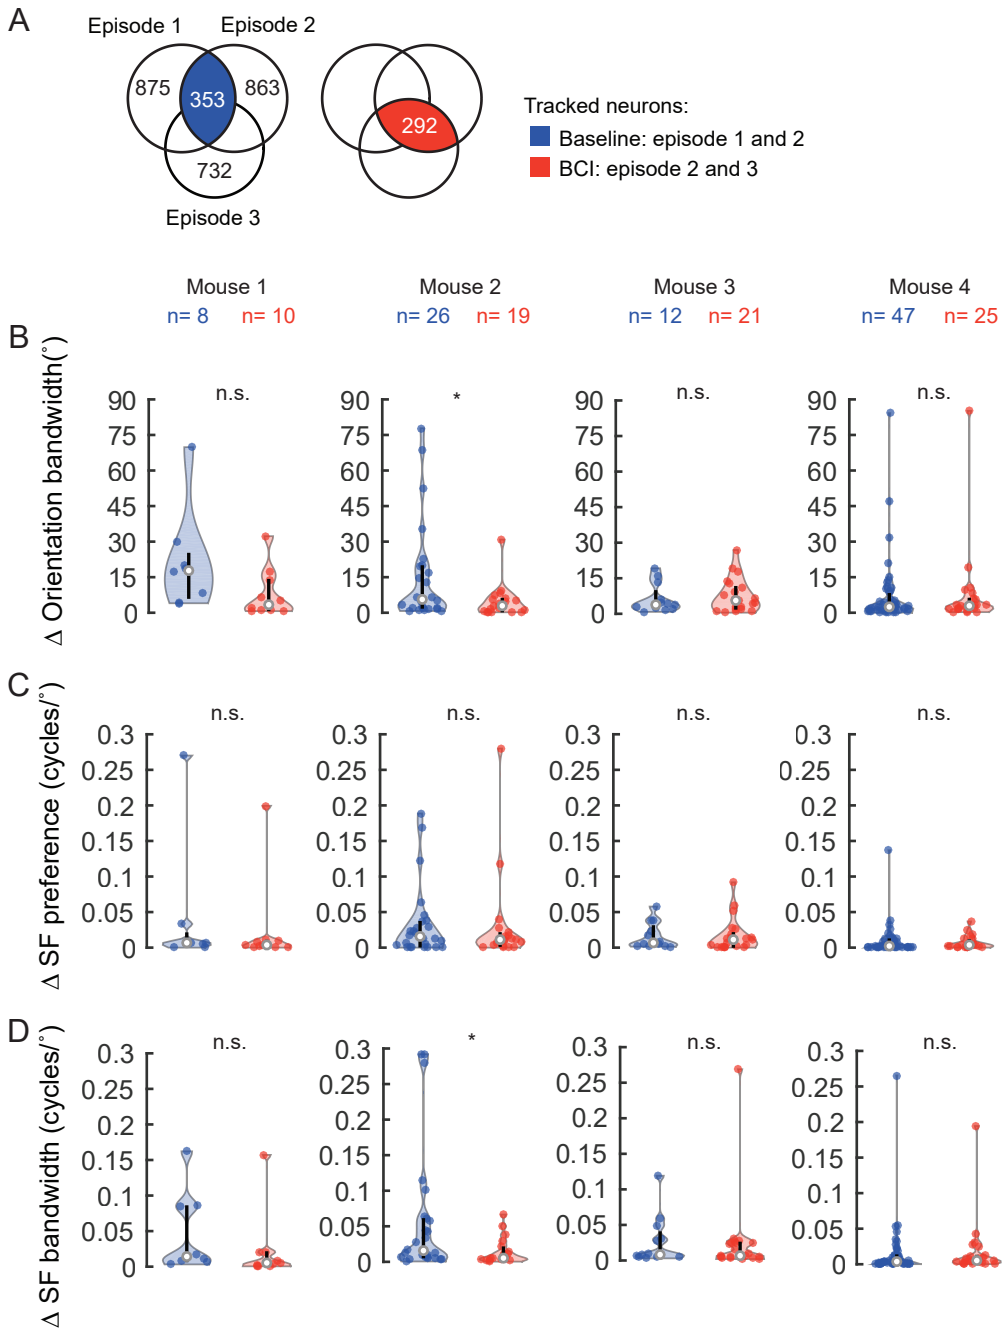

Supplement Figure 5. Orientation and spatial frequency tuning features remain stable after the acquisition of BCI control.

(A) Venn diagrams, re-plotted from Fig. 4A.

(B-D) Stability of tuning features in the baseline condition (blue) and BCI condition (red) as indicated. Labels as in Fig. 3D, 4D. Number of neurons per animal indicated (same as Fig. 4D). Distributions for the three features were overlapping in all 4 mice. Wilcoxon rank-sum test p values (mouse 1-4, respectively): *Orientation bandwidth*, p = 0.055, p = 0.042, p = 0.667, p = 0.970. *Spatial frequency preference*, p = 0.573, p = 0.573, p = 0.587, p = 0.940. *Spatial frequency bandwidth*, p = 0.122, p = 0.038, p = 0.270, p = 0.361.

Supplement Figure 6

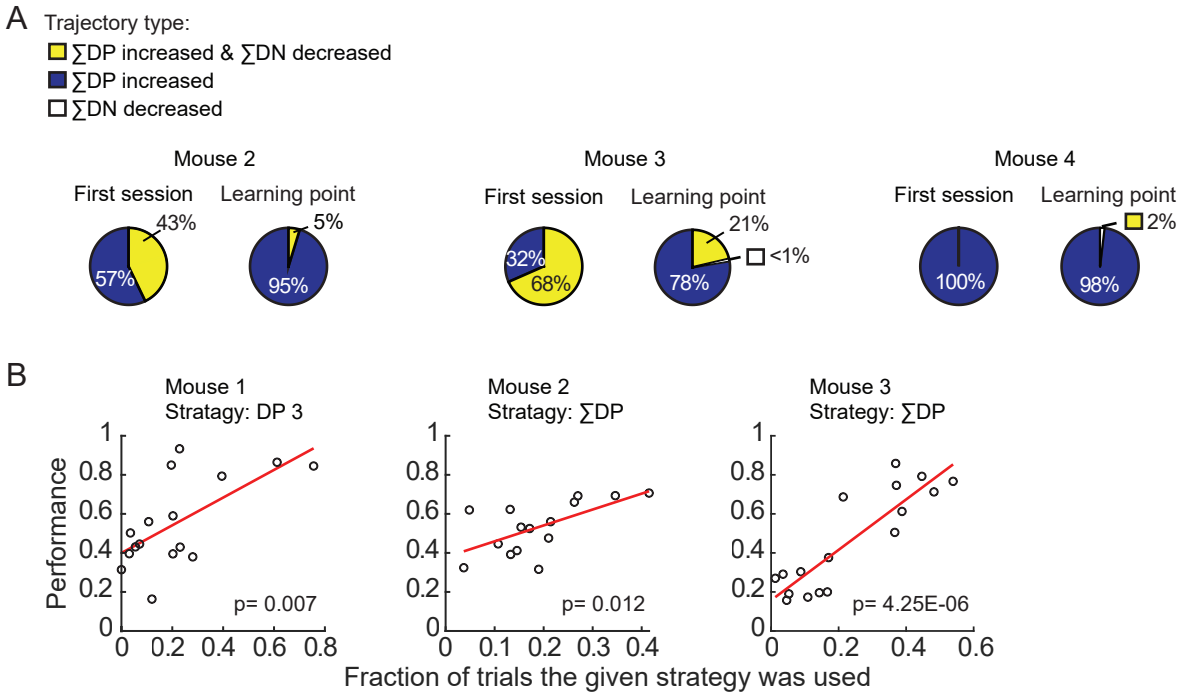

Supplement Figure 6. Performance level was associated with a specific strategy in 3 of 4 mice.

(A) The trajectory type used to reach the target was largely unchanged by training. Trajectory type for Mouse 1 is shown in Figure 5.

(B) In 3 of 4 mice, one of the four possible strategies that employed direct positive neurons significantly correlated with performance. Regression analysis was used to assess significance, only the strategies with significant correlation are shown.