Report

Current Biology

Brain-Computer Interface with Inhibitory Neurons Reveals Subtype-Specific Strategies

Highlights

- Mice can learn to modulate activity difference of two inhibitory neurons
- Mice employ subtype-specific strategies to modulate each neuron
- When PV neurons are targeted, activity of negative target decreases
- When SOM and VIP neurons are targeted, activity of positive target increases

Authors

Akinori Mitani, Mingyuan Dong, Takaki Komiyama

Correspondence

tkomiyama@ucsd.edu

In Brief

Mitani et al. demonstrate that mice can modulate the activity difference of a pair of PV-, SOM-, and VIP-expressing inhibitory neurons with subtype-specific strategies in a brain-computer interface task using two-photon imaging of inhibitory neurons of each subtype expressing GCaMP6f.

Mitani et al., 2018, Current Biology *28*, 77–83
January 8, 2018
https://doi.org/10.1016/j.cub.2017.11.035

CellPress

Current Biology
Report

Brain-Computer Interface with Inhibitory Neurons Reveals Subtype-Specific Strategies

Akinori Mitani,[1] Mingyuan Dong,[1] and Takaki Komiyama[1,2,*]
[1]Neurobiology Section, Center for Neural Circuits and Behavior, and Department of Neurosciences, University of California, San Diego, La Jolla, CA 92093, USA
[2]Lead Contact
*Correspondence: tkomiyama@ucsd.edu
https://doi.org/10.1016/j.cub.2017.11.035

SUMMARY

Brain-computer interfaces have seen an increase in popularity due to their potential for direct neuroprosthetic applications for amputees and disabled individuals. Supporting this promise, animals—including humans—can learn even arbitrary mapping between the activity of cortical neurons and movement of prosthetic devices [1–4]. However, the performance of neuroprosthetic device control has been nowhere near that of limb control in healthy individuals, presenting a dire need to improve the performance. One potential limitation is the fact that previous work has not distinguished diverse cell types in the neocortex, even though different cell types possess distinct functions in cortical computations [5–7] and likely distinct capacities to control brain-computer interfaces. Here, we made a first step in addressing this issue by tracking the plastic changes of three major types of cortical inhibitory neurons (INs) during a neuron-pair operant conditioning task using two-photon imaging of IN subtypes expressing GCaMP6f. Mice were rewarded when the activity of the positive target neuron (N+) exceeded that of the negative target neuron (N−) beyond a set threshold. Mice improved performance with all subtypes, but the strategies were subtype specific. When parvalbumin (PV)-expressing INs were targeted, the activity of N− decreased. However, targeting of somatostatin (SOM)- and vasoactive intestinal peptide (VIP)-expressing INs led to an increase of the N+ activity. These results demonstrate that INs can be individually modulated in a subtype-specific manner and highlight the versatility of neural circuits in adapting to new demands by using cell-type-specific strategies.

RESULTS AND DISCUSSION

Water-restricted mice expressing GCaMP6f in parvalbumin (PV)-, somatostatin (SOM)-, or vasoactive intestinal peptide (VIP)-expressing inhibitory neurons (INs) were trained in a neuron-pair operant conditioning task with two-photon calcium imaging (modified from [8]). Briefly, two neurons in layer 2/3 of the primary motor cortex were randomly selected out of those that met predetermined activity criteria (STAR Methods) and designated as positive target neuron (N+) and negative target neuron (N−). In some of the imaging fields (14/36), there were only two labeled neurons that met the activity criteria. The distance between N+ and N− was 191 ± 119 μm (mean ± SD; n = 36 pairs). The calcium signal of the targeted neurons was not saturated (Figures S1A–S1C). Mice were rewarded when the calcium activity of N+ exceeded that of N− by a set threshold (Figure 1A). The reward contingency based on the difference between N+ and N− activity ensured that mice could not solve the task by simply activating all neurons in the area simultaneously. During the trial, the difference of the calcium signal of the two neurons was transformed to create a dynamically frequency-modulated auditory feedback. After each reward, the activity of the targeted neurons had to return to baseline, which resumed the auditory feedback and initiated the next trial. The same neurons were targeted over 4–6 sessions, one session per day, with the same reward contingency. We note that a previous study with a similar approach that targeted neurons of unidentified cell types showed that auditory feedback was essential for the learning of the task [8]. However, in the current study, we did not explicitly test the necessity of the auditory feedback.

For each of the three major IN types, mice were able to improve the performance over sessions, significantly increasing the reward frequency (PV-INs, $p = 0.011$, n = 48 sessions, 10 imaging fields in 5 mice, using fitlme in MATLAB hereafter unless otherwise stated; SOM-INs, $p = 0.001$, 50 sessions, 10 imaging fields in 5 mice; VIP-INs, $p = 0.017$, 82 sessions, 16 imaging fields in 8 mice; Figures 1B–1D). Immunostaining showed a high degree of overlap between GCaMP6f-expressing and PV-expressing neurons in PV animals and little overlap between GCaMP6f-expressing and PV-expressing neurons in SOM and VIP animals (Figures 1E–1G). In a subset of the animals, we imaged the same neurons for an additional 1–3 contingency degradation sessions, in which rewards were provided without regard to the activity of the targeted neurons. In these sessions, the targeted neurons reached the reward threshold significantly less frequently (the difference between the contingency degradation sessions and the last two training sessions, 2.49 ± 0.64/min [estimate ± SE; $p < 0.001$; n = 12 and 14

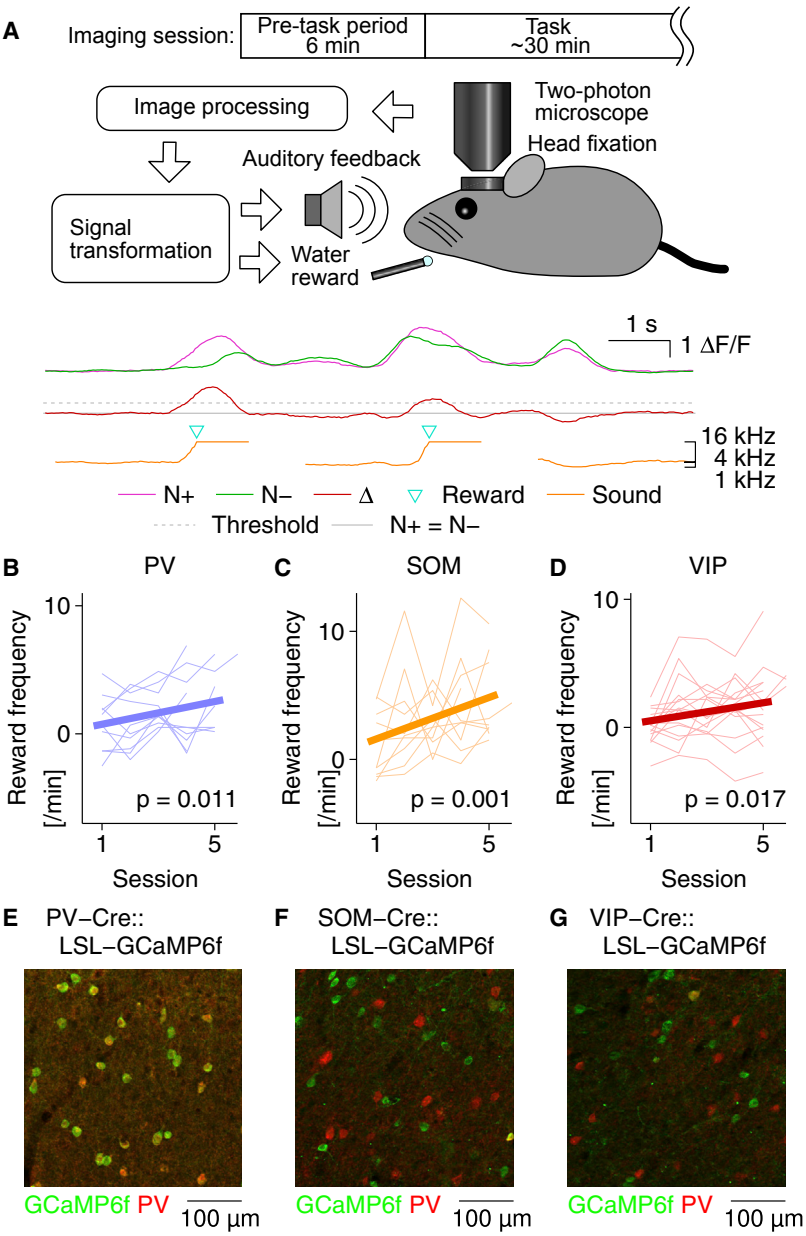

Figure 1. Mice Improved Performance in a Neuron-Pair Operant Conditioning Task with IN Subtypes

(A) Schematic of task. Δ, N+-N−.

(B–D) Reward frequency, adjusted for chance-level performance during pre-task period in each session, across 4–6 sessions in PV (B; 0.44 ± 0.17 [slope estimate ± SE]; n = 48 sessions; 10 fields in 5 mice; using fitlme in MATLAB), SOM (C; 0.81 ± 0.24; 50 sessions; 10 fields in 5 mice), and VIP mice (D; 0.35 ± 0.14; 82 sessions; 16 fields in 8 mice). Thick lines, linear fit; thin lines, each imaging field.

(E) Example PV (red) and GCaMP6f (green) labeling in PV animals. In total, among 309 GCaMP6f-positive cells and 289 PV-antibody-stained cells, 279 were double positive.

(F) Same as (E) in SOM animals. Among 196 GCaMP6f-positive cells and 176 PV-antibody-stained cells, only 7 were double positive.

(G) Same as (E) in VIP animals. Among 164 GCaMP6f-positive cells and 202 PV-antibody-stained cells, only 4 were double positive.

See also Figure S1.

sessions; 7 imaging fields in 4 mice], 3.07 ± 0.71/min [p < 0.001; n = 20 and 20 sessions; 10 imaging fields in 5 mice], and 1.08 ± 0.38/min [p = 0.006; n = 28 and 30 sessions; 15 imaging fields in 7 mice] in PV, SOM, and VIP animals, respectively). These results indicate that mice are indeed able to modulate the activity of IN subtypes.

We considered three possible strategies by which mice could achieve an increase in reward frequency (Figure 2A). For example, the reward frequency could increase with an increase in the activity frequency of N+ (Figure 2A; "N+ increase"). Alternatively, a decrease in the activity of N− (Figure 2A; "N− decrease") or decoupling of activity between N+ and N− (Figure 2A; "decoupling") can improve the reward frequency by increasing the chance that activity in N+ leads to a reward. In the following analysis, we focused on the frequency of calcium events because the amplitude of the events did not significantly change across sessions (Figures S1D–S1F).

When PV-INs were targeted (Figure 2B), the frequency of N+ calcium events did not change significantly, arguing against the N+ increase strategy (Figure 2C). In contrast, the frequency of N− calcium events decreased (p = 0.002), supporting the N− decrease strategy (Figure 2D). The slopes of the changes of N+ activity (Figure 2C) and N− activity (Figure 2D) were significantly different (p = 0.001), indicating the specificity of the decrease of N− activity. Accordingly, the frequency of the N+ calcium events that did not lead to a reward decreased (p = 0.025; Figure 2E). The frequency of N− events during the time periods when N+ was inactive also significantly decreased over sessions (p = 0.011; Figure 2F). This result argues against the decoupling strategy, which would predict that N− events that originally coincided with N+ activity would move into the periods of N+ inactivity. Furthermore, the correlation coefficient between N+ and N− activity during task period did not decrease across sessions (slope estimate ± SE = 0.0145 ± 0.0089/session; p = 0.112). To test whether the changes of neural activity contribute to the improvement in task performance, we conducted mediation analysis [9]. In this analysis, we found that N+ activity and N− activity were positively and negatively correlated with task performance, respectively (coefficients of a linear model: 0.57 ± 0.06 [p < 0.001] and −0.22 ± 0.06 [p < 0.001; estimate ± SE]), and there was a significant mediation effect with N− activity decrease (p < 0.001; STAR Methods) and not with N+ activity (p = 0.233). These results indicate that, when PV-INs were targeted, mice specifically decreased the activity of N− while maintaining N+ activity.

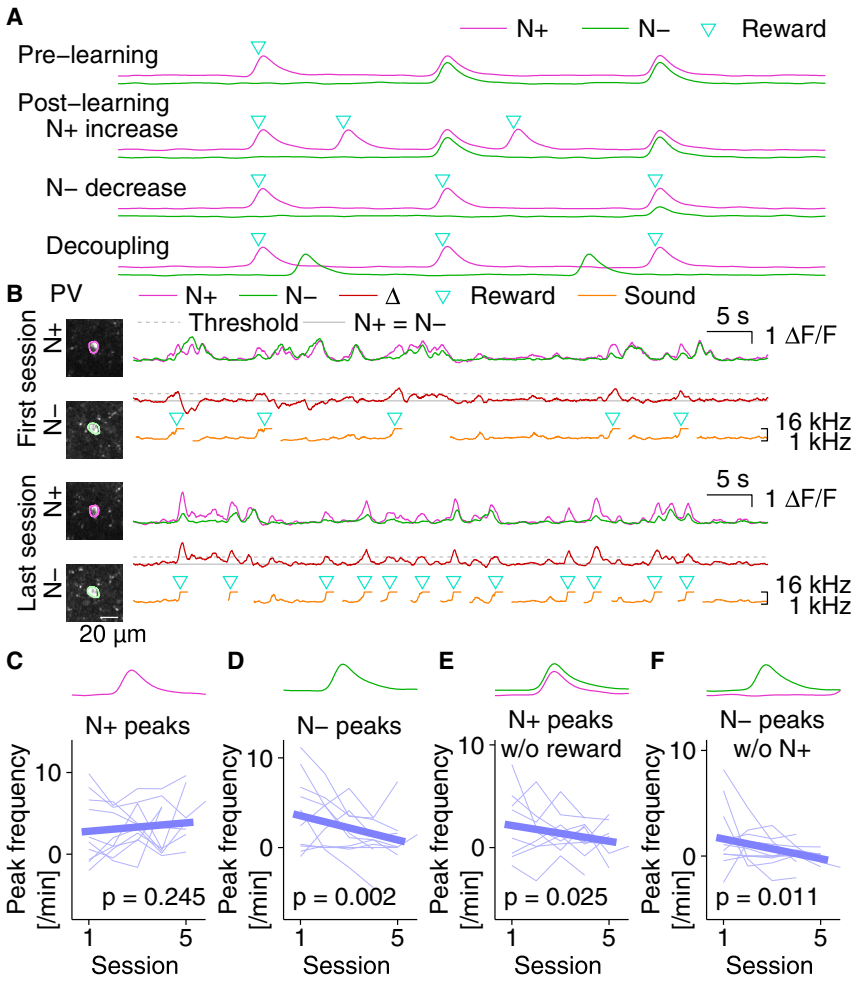

Figure 2. Neural Activity Modulation during PV-IN Targeting

(A) Schematic of three potential strategies to achieve the task.

(B) Example activity traces of N+ and N−, activity difference (Δ), auditory feedback, and reward timing in the first (top) and the last (bottom) session of a PV mouse.

(C–F) Daily changes of the frequency of N+ peaks (C; 0.26 ± 0.2 [slope estimate ± SE]; n = 48 sessions; 10 fields in 5 mice; using fitlme in MATLAB), N− peaks (D; −0.68 ± 0.22), N+ peaks that were not associated with rewards (E; −0.39 ± 0.17), and N− peaks during the periods of N+ inactivity (F; −0.47 ± 0.18) during PV-IN targeting, relative to pre-task period.

The task improvement in mice with SOM-IN targeting (Figure 3A) involved a different strategy. In contrast to PV-INs, N+ event frequency increased in the later sessions in SOM-INs ($p = 0.003$; Figure 3B). This suggests that the N+ increase strategy was utilized to perform the task. Conversely, the frequency of N− events did not change (Figure 3C), nor did the frequency of N+ events that were not associated with rewards (Figure 3D), arguing against the N− decrease hypothesis. The slopes of the activity changes of N+ (Figure 3B) and N− (Figure 3C) were significantly different ($p = 0.010$). Furthermore, the frequency of N− events during the periods of N+ inactivity did not increase (Figure 3E) and the correlation coefficient between N+ and N− activity during task period did not significantly decrease across sessions (slope estimate ± SE = −0.0009 ± 0.0119/session; $p = 0.939$), discounting the decoupling strategy. In the mediation analysis, N+ and N− activities were positively and negatively correlated with task performance, respectively (coefficients of a linear model: 0.60 ± 0.03 [$p < 0.001$] and −0.06 ± 0.02 [$p = 0.018$; estimate ± SE]), and the mediation effect through N+ activity was significant ($p = 0.003$) and not through N− activity ($p = 0.082$). We conclude that mice improved the task performance with SOM-INs primarily by activating specifically the N+ neuron.

Next, we investigated VIP-INs (Figure 3F). As with SOM-INs, in VIP mice, the event frequency of N+ significantly increased ($p = 0.044$; Figure 3G). Neither the frequency of N− events nor the N+ event frequency not associated with reward changed significantly (Figures 3H and 3I), arguing against the N− decrease strategy, although the difference in the slopes of activity changes of N+ (Figure 3G) and N− (Figure 3H) did not reach statistical significance ($p = 0.101$). The N− event frequency within the periods of N+ inactivity did not change (Figure 3J), and the correlation coefficient between N+ and N− activity during task period did not change (slope estimate ± SE = −0.0103 ± 0.0091/session; $p = 0.261$), excluding the decoupling strategy. N+ and N− activity were positively and negatively correlated with task performance, respectively (coefficients of a linear model: 0.70 ± 0.03 [$p < 0.001$] and −0.10 ± 0.04 [$p = 0.011$; estimate ± SE]), and the mediation effect through N+ activity was significant ($p = 0.040$). These data demonstrate that VIP mice improved the task performance by increasing N+ activity.

To test whether the difference between cell types was significant, we examined whether cell type had a significant effect on the slope of the linear model (STAR Methods; Figures 4A–4D). The reward frequency increase was not different among three cell types. The frequency increase of N+ peaks was significantly larger in SOM neurons than in PV and VIP neurons. The decrease of N− peak frequency was greater in PV neurons than in VIP neurons, and the decrease of the frequency of N+ peaks without rewards was specific to PV.

For SOM and VIP mice, we further investigated whether N+/N− activity changed differently from the activity of non-target neurons, which met the same activity criteria. We did not perform this analysis for PV mice because 6 out of 10 imaging fields did not have any non-target neuron that met the activity criteria and the other 4 fields only had 1 non-target neuron. The result shows that N+ activity increase in SOM animals was specific to N+ and significantly greater than the non-target neurons (slope difference from non-target

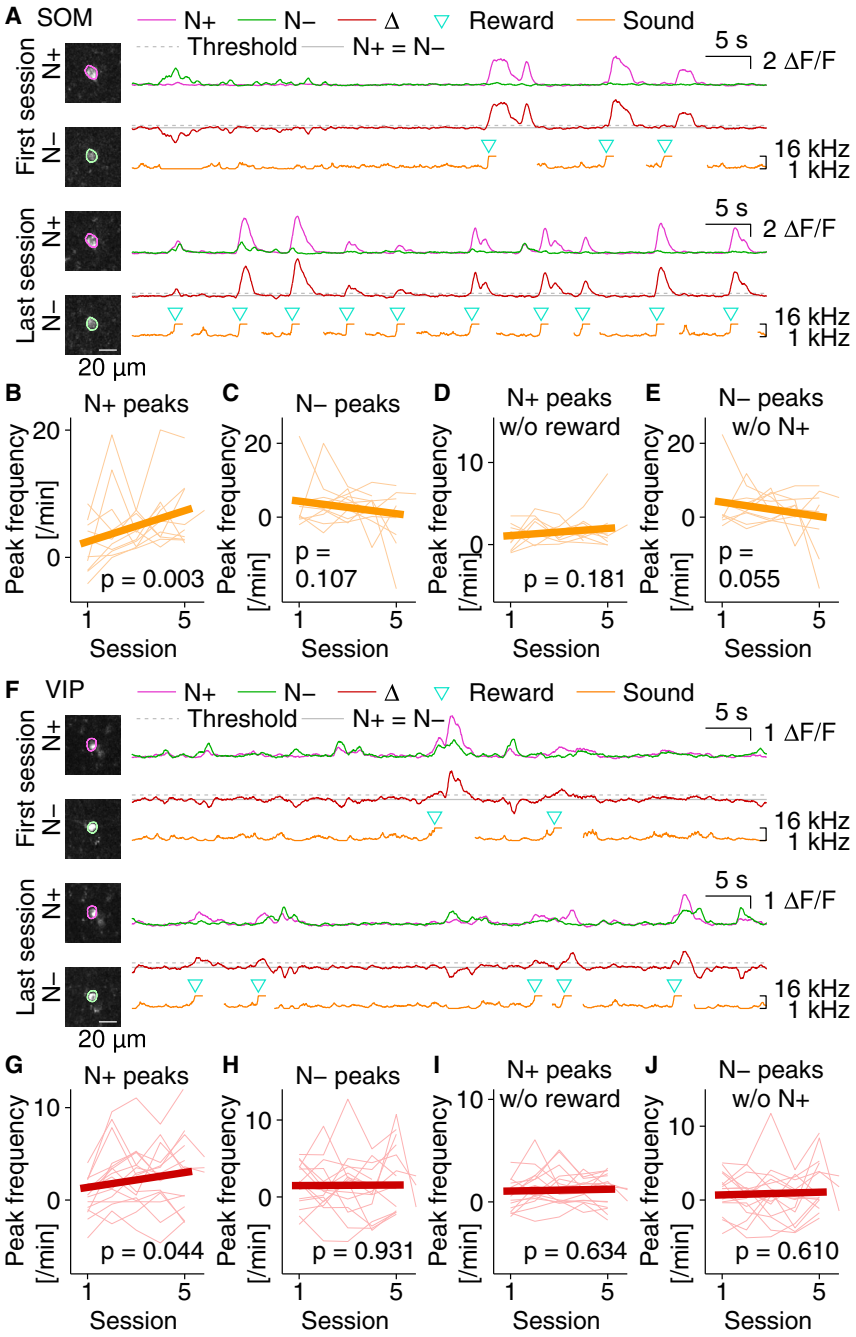

Figure 3. Neural Activity Modulation during SOM-IN and VIP-IN Targeting

(A) Example activity traces of N+ and N–, activity difference (Δ), auditory feedback, and reward timing in the first (top) and the last (bottom) session of a SOM mouse.

(B–E) Daily changes of the frequency of N+ peaks (B; 1.21 ± 0.39 [slope estimate ± SE]; n = 50 sessions; 10 fields in 5 mice; using fitlme in MATLAB), N– peaks (C; 0.84 ± 0.51), N+ peaks that were not associated with rewards (D; 0.22 ± 0.16), and N– peaks during the periods of N+ inactivity (E; −0.97 ± 0.50) during SOM-IN targeting, relative to pre-task period.

(F) Same as (A) for a VIP mouse.

(G–J) Same as (B)–(E) for VIP-INs (n = 82 sessions; 16 fields in 8 mice). Slope estimates ± SE are 0.40 ± 0.20 (G), 0.02 ± 0.21 (H), 0.04 ± 0.09 (I), and 0.09 ± 0.18 (J).

neurons: 1.81 ± 0.73 [p = 0.014; SOM; N+; Figure 4E], −0.68 ± 0.77 [p = 0.376; SOM; N–; Figure 4F], 0.43 ± 0.41 [p = 0.304; VIP; N+; Figure 4E] and −0.12 ± 0.42 [p = 0.770; VIP; N–; Figure 4F; estimate ± SE]). In addition, we simulated reward frequency increase as if each non-target neuron were used in the task as either N+ or N–, and the actual target neuron was used for the other target (for example, in a simulation in which we used a non-target neuron as N+, reward frequency was simulated using that non-target neuron as N+ and the actual N–). Figure 4G shows that the simulated reward frequency increase with non-target neurons used as N+ is significantly lower than the actual reward frequency increase in SOM neurons, but not in VIP neurons. The reward frequency increase was not affected in either SOM or VIP animals when non-target neurons were used as N– (Figure 4H). The results show that the activity increase in SOM animals was specific to N+, leading to improved task performance.

To our knowledge, this is the first study to test the plasticity of individual neurons of molecularly identified cell types in a brain-computer interface task. We demonstrate that cell type has a profound impact on the way by which performance improvement is achieved. For example, the activity of a SOM-IN could be increased without activating a second SOM-IN, similarly to a previous study that did not identify the cell type of the targeted neurons (and thus most of the targeted neurons were presumably excitatory) [8]. However, we found no evidence that the activity of a PV-IN could be increased without also activating a second PV-IN. Instead, the activity of a PV-IN could be reduced without inactivating a second PV-IN. The performance of brain-machine interfaces may improve in the future if such cell-type-specific constraints on plasticity are considered [10].

The differences in baseline activity levels may have partially contributed to the difference in strategies among subtypes. For example, if the baseline activity level of PV neurons is higher, it might be more difficult to increase N+ activity in PV-INs than in other subtypes. In addition, due to different calcium buffering in each cell, the relationship between spikes and GCaMP6f signals may be different from cell to cell, leaving the possibility for PV neurons to require more spikes to cause a calcium event.

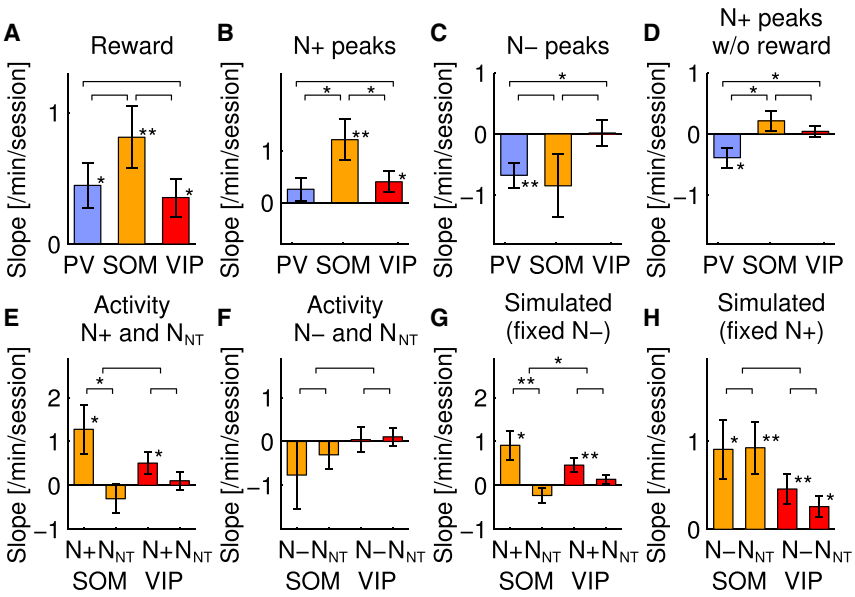

Figure 4. Difference between Cell Types in Neural Activity Modulation

(A–D) Difference between cell types in the slopes of the linear model. The frequency of rewards (A), N+ peaks (B), N− peaks (C), and N+ peaks that were not associated with rewards (D) is shown.

(E and F) Difference between activity increase of target neurons (N+ [E] and N− [F]) and that of non-target neurons (N_{NT}).

(G and H) Difference between actual reward frequency increase and simulated reward frequency increase if activity of non-target neurons was used instead of N+ (G) or N− (H). Bar plot and error bars indicate slope estimates ± SE (*$p < 0.05$; **$p < 0.01$). * and ** between bars indicate a significant interaction term of a linear model (*$p < 0.05$; **$p < 0.01$). In (E)–(H), only imaging fields that had non-target neurons were included in the analysis. See also Figure S2.

Alternatively, the cell-type-specific strategies may be partially explained by the differences in the levels of activity correlation within each cell type (pairwise correlation coefficients of pre-task activity between candidate neurons were PV, 0.61 ± 0.84 [n_{pair} = 18]; SOM, 0.14 ± 0.03 [n_{pair} = 40]; and VIP, 0.31 ± 0.01 [n_{pair} = 394]; mean ± SEM; $p < 0.001$ for all pairwise comparisons after removing the effect of event rate; Figure S2). Nevertheless, we argue that the difference in correlation is an important reflection of their intrinsic properties.

The ability to improve task performance with PV-INs is particularly striking, given the high activity correlation between PV-INs. It has been shown that neural feedback tasks based on the difference of two neural ensembles is harder to learn if the activity of the two ensembles is correlated [8]. However, the animals could perform the task with a specific strategy.

What could be the potential mechanisms underlying cell-type specificity of strategies? At the cellular level, the performance of this task could be mediated by specific plasticity of N+ and/or N− neurons, such as plasticity of intrinsic excitability and synaptic plasticity of inhibitory and excitatory synapses onto these neurons. There are likely differences among cell types in their ability for these plasticity mechanisms (reviewed in [11]). For example, repetitive correlated spiking induced spike-timing-dependent plasticity (STDP) in low-threshold spiking (LTS) interneurons (putative non-PV INs) [12], whereas it induced long-term depression in fast spiking interneurons (putative PV-INs) [12, 13]. If excitatory synapses onto PV neurons are less likely to be potentiated when the neurons are targeted in the operant conditioning task, it can explain why the N− decrease strategy was employed with PV-INs. However, it has been shown that long-term potentiation can be induced in PV-INs with theta burst stimulation [14], and synaptic plasticity is sensitive to the neuromodulatory state of the circuit [15] and behavioral context [16, 17]. Therefore, it remains unclear how the capacities for synaptic plasticity may differ across cell types in the intact brain during learning.

Another potential mechanism for the improved performance in the current task is by modulating the activity of neurons presynaptic to the targeted neurons. Different subtypes of inhibitory neurons receive inputs from different populations of excitatory neurons. PV-INs receive dense inputs from nearby excitatory neurons [18], which suggests that nearby PV-INs share similar excitatory inputs and thus a specific increase of excitatory inputs to one PV-IN may be difficult. On the other hand, SOM- and VIP-INs receive excitatory inputs from largely non-overlapping populations [19]. Differences in the strategy for the operant conditioning task may originate from the differences in the characteristics of the neurons providing excitatory inputs to the target neurons.

In addition, excitatory and inhibitory neurons form highly interconnected networks. In general, it is thought that VIP-INs inhibit SOM-INs; SOM-INs inhibit excitatory, PV-, and VIP-INs; PV-INs inhibit PV-INs and excitatory neurons [6]; and excitatory neurons project to all four types. These connectivity patterns provide many possible pathways that could mediate the plasticity observed in the current study. For example, for SOM-INs, SOM → PV → excitatory → SOM provides a potential positive feedback loop, possibly underlying the increase of N+ activity in SOM-INs. Furthermore, the SOM → PV inhibition could underlie the decrease of N− activity in PV-INs. Future experiments are required to test these specific possibilities.

Finally, the activity of the inhibitory neurons can be associated with movements and sensory stimulus in a subtype-specific manner. A study showed that monkeys perform a brain-machine interface task by exploring and exploiting neural patterns associated with natural movements [20]. Through our visual observations, we did not identify overt behavioral strategies during task performance, similar to a previous study [8]. Nevertheless, examining the relationship between existing neural activity patterns and behavioral variables (“intrinsic neural manifold” and “intrinsic behavioral manifold”)

[21] before training and how they change through training in each subtype will be of future interest. Furthermore, studies using neural feedback task with two-photon calcium imaging have reported that auditory [8] or artificial sensory [22] feedback was necessary for the successful learning of the task, whereas similar learning only with rewards as a feedback has also been reported [23]. Future experiments can be aimed at examining whether sensory feedback was necessary to learn to modulate inhibitory neurons and, if so, how the dependency is different among cell types. Lastly, we note that our contingency degradation experiments suggest that the behavioral performance was goal directed. However, the lack of auditory feedback in the degradation experiments leaves room for other interpretations, such as that the mice might have been in a completely different behavioral state without the feedback. It will be of future interest to investigate the relationships between targeted cell types and behavioral strategies, dependence on sensory feedback, and whether the behavioral performance is goal directed.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- CONTACT FOR REAGENT AND RESOURCE SHARING
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
  - Animals
- METHOD DETAILS
  - Immunostaining and cell counting
  - Surgery
  - Imaging
  - Behavioral task
- QUANTIFICATION AND STATISTICAL ANALYSES
  - Data analysis
  - Linear models
  - Statistics
- DATA AND SOFTWARE AVAILABILITY

SUPPLEMENTAL INFORMATION

Supplemental Information includes four figures and can be found with this article online at https://doi.org/10.1016/j.cub.2017.11.035.

ACKNOWLEDGMENTS

We thank A.N. Kim, K. O'Neil, L. Hall, T. Loveland, and O. Arroyo for technical assistance; the Janelia GENIE Project for GCaMP6f; and members of the Komiyama lab, especially R. Hattori, N. Hedrick, E. Hwang, and H. Liu, for comments and discussions. This work was supported by grants from NIH (R01 NS091010A, R01 EY025349, R01 DC014690, U01 NS094342, and P30EY022589), Pew Charitable Trusts, David and Lucile Packard Foundation, McKnight Foundation, New York Stem Cell Foundation, Kavli Institute for Brain and Mind, and NSF (1734940) to T.K. A.M. is supported by the Nakajima Foundation.

AUTHOR CONTRIBUTIONS

T.K. and A.M. conceived the project and wrote the paper. Immunostaining was performed by M.D. and A.M. All other experiments were performed by A.M. and analyzed by A.M. and T.K.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: July 7, 2017
Revised: October 10, 2017
Accepted: November 15, 2017
Published: December 14, 2017

REFERENCES

1. Moritz, C.T., Perlmutter, S.I., and Fetz, E.E. (2008). Direct control of paralysed muscles by cortical neurons. Nature *456*, 639–642.
2. Jarosiewicz, B., Chase, S.M., Fraser, G.W., Velliste, M., Kass, R.E., and Schwartz, A.B. (2008). Functional network reorganization during learning in a brain-computer interface paradigm. Proc. Natl. Acad. Sci. USA *105*, 19486–19491.
3. Ganguly, K., and Carmena, J.M. (2009). Emergence of a stable cortical map for neuroprosthetic control. PLoS Biol. *7*, e1000153.
4. Sadtler, P.T., Quick, K.M., Golub, M.D., Chase, S.M., Ryu, S.I., Tyler-Kabara, E.C., Yu, B.M., and Batista, A.P. (2014). Neural constraints on learning. Nature *512*, 423–426.
5. Chen, S.X., Kim, A.N., Peters, A.J., and Komiyama, T. (2015). Subtype-specific plasticity of inhibitory circuits in motor cortex during motor learning. Nat. Neurosci. *18*, 1109–1115.
6. Pfeffer, C.K., Xue, M., He, M., Huang, Z.J., and Scanziani, M. (2013). Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons. Nat. Neurosci. *16*, 1068–1076.
7. Pi, H.-J., Hangya, B., Kvitsiani, D., Sanders, J.I., Huang, Z.J., and Kepecs, A. (2013). Cortical interneurons that specialize in disinhibitory control. Nature *503*, 521–524.
8. Clancy, K.B., Koralek, A.C., Costa, R.M., Feldman, D.E., and Carmena, J.M. (2014). Volitional modulation of optically recorded calcium signals during neuroprosthetic learning. Nat. Neurosci. *17*, 807–809.
9. Falk, C.F., and Biesanz, J.C. (2016). Two cross-platform programs for inferences and interval estimation about indirect effects in mediational models. SAGE Open *6*, 2158244015625445.
10. Shenoy, K.V., and Carmena, J.M. (2014). Combining decoder design and neural adaptation in brain-machine interfaces. Neuron *84*, 665–680.
11. Kullmann, D.M., Moreau, A.W., Bakiri, Y., and Nicholson, E. (2012). Plasticity of inhibition. Neuron *75*, 951–962.
12. Yavorska, I., and Wehr, M. (2016). Somatostatin-expressing inhibitory interneurons in cortical circuits. Front. Neural Circuits *10*, 76.
13. Lu, J.T., Li, C.Y., Zhao, J.-P., Poo, M.M., and Zhang, X.H. (2007). Spike-timing-dependent plasticity of neocortical excitatory synapses on inhibitory interneurons depends on target cell type. J. Neurosci. *27*, 9711–9720.
14. Sarihi, A., Jiang, B., Komaki, A., Sohya, K., Yanagawa, Y., and Tsumoto, T. (2008). Metabotropic glutamate receptor type 5-dependent long-term potentiation of excitatory synapses on fast-spiking GABAergic neurons in mouse visual cortex. J. Neurosci. *28*, 1224–1235.
15. Huang, S., Huganir, R.L., and Kirkwood, A. (2013). Adrenergic gating of Hebbian spike-timing-dependent plasticity in cortical interneurons. J. Neurosci. *33*, 13171–13178.
16. Donato, F., Rompani, S.B., and Caroni, P. (2013). Parvalbumin-expressing basket-cell network plasticity induced by experience regulates adult learning. Nature *504*, 272–276.
17. Kuhlman, S.J., Olivas, N.D., Tring, E., Ikrar, T., Xu, X., and Trachtenberg, J.T. (2013). A disinhibitory microcircuit initiates critical-period plasticity in the visual cortex. Nature *501*, 543–546.
18. Hofer, S.B., Ko, H., Pichler, B., Vogelstein, J., Ros, H., Zeng, H., Lein, E., Lesica, N.A., and Mrsic-Flogel, T.D. (2011). Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex. Nat. Neurosci. *14*, 1045–1052.

19. Karnani, M.M., Jackson, J., Ayzenshtat, I., Tucciarone, J., Manoocheri, K., Snider, W.G., and Yuste, R. (2016). Cooperative subnetworks of molecularly similar interneurons in mouse neocortex. Neuron *90*, 86–100.

20. Hwang, E.J., Bailey, P.M., and Andersen, R.A. (2013). Volitional control of neural activity relies on the natural motor repertoire. Curr. Biol. *23*, 353–361.

21. Jazayeri, M., and Afraz, A. (2017). Navigating the neural space in search of the neural code. Neuron *93*, 1003–1014.

22. Prsa, M., Galiñanes, G.L., and Huber, D. (2017). Rapid integration of artificial sensory feedback during operant conditioning of motor cortex neurons. Neuron *93*, 929–939.e6.

23. Hira, R., Ohkubo, F., Masamizu, Y., Ohkura, M., Nakai, J., Okada, T., and Matsuzaki, M. (2014). Reward-timing-dependent bidirectional modulation of cortical microcircuits during optical single-neuron operant conditioning. Nat. Commun. *5*, 5551.

24. Hippenmeyer, S., Vrieseling, E., Sigrist, M., Portmann, T., Laengle, C., Ladle, D.R., and Arber, S. (2005). A developmental switch in the response of DRG neurons to ETS transcription factor signaling. PLoS Biol *3*, e159.

25. Madisen, L., Garner, A.R., Shimaoka, D., Chuong, A.S., Klapoetke, N.C., Li, L., van der Bourg, A., Niino, Y., Egolf, L., Monetti, C., et al. (2015). Transgenic mice for intersectional targeting of neural sensors and effectors with high specificity and performance. Neuron *85*, 942–958.

26. Taniguchi, H., He, M., Wu, P., Kim, S., Paik, R., Sugino, K., Kvitsiani, D., Fu, Y., Lu, J., Lin, Y., et al. (2011). A resource of Cre driver lines for genetic targeting of GABAergic neurons in cerebral cortex. Neuron *71*, 995–1013.

27. Peters, A.J., Chen, S.X., and Komiyama, T. (2014). Emergence of reproducible spatiotemporal activity during motor learning. Nature *510*, 263–267.

28. Evangelidis, G.D., and Psarakis, E.Z. (2008). Parametric image alignment using enhanced correlation coefficient maximization. IEEE Trans. Pattern Anal. Mach. Intell. *30*, 1858–1865.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Antibodies		
Rabbit anti-PV	Abcam	Cat#ab11427; RRID: AB_298032
Chicken anti-GFP	Aves Labs	Cat#GFP-1020; RRID: AB_10000240
Goat anti-chicken DyLight 488	Thermo Fisher Scientific	Cat#SA5-10070; RRID: AB_2556650
Donkey anti-rabbit Alexa 594	Thermo Fisher Scientific	Cat#A-21207; RRID: AB_141637
Experimental Models: Organisms/Strains		
PV-Cre	The Jackson Laboratory	RRID: IMSR_JAX:017320
SOM-Cre	The Jackson Laboratory	RRID: IMSR_JAX:013044
VIP-Cre	The Jackson Laboratory	RRID: IMSR_JAX:010908
lsl-GCaMP6F	The Jackson Laboratory	RRID: IMSR_JAX:024105
VIP-Cre::ZtTA::TITL-GCaMP6f	The Jackson Laboratory	RRID: IMSR_JAX:024107
Software and Algorithms		
ScanImage 4	Vidrio Technologies	RRID: SCR_014307
MATLAB R2014a	The MathWorks	RRID: SCR_001622

CONTACT FOR REAGENT AND RESOURCE SHARING

Further information and requests for resources and reagents should be directed to and will be fulfilled by the Lead Contact, Takaki Komiyama (tkomiyama@ucsd.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Animals

All procedures were in accordance with protocols approved by UCSD Institutional Animal Care and Use Committee and guidelines of the US National Institutes of Health. All animals before water restriction were group housed, and during water restriction they were singly housed or group housed when all littermates were under water restriction. They were housed in disposable plastic cages with standard bedding in a room on a reversed light cycle (12h/12h). Experiments were typically performed during the dark period. With mice acquired from Jackson Laboratories, we generated double transgenic mice, PV-Cre [JAX:017320] [24]::lsl-GCaMP6f [JAX:024105] [25], SOM-Cre [JAX:013044] [26]::lsl-GCaMP6f, VIP-Cre [JAX:010908] [26]::lsl-GCaMP6f, and triple transgenic mice VIP-Cre::ZtTA::TITL-GCaMP6f [JAX:024107] [25] to express the calcium indicator GCaMP6f in a specific subtype of inhibitory neurons. Each line was maintained in the original background and crossed to generate the double or triple transgenic mice for experiments. Double transgenic and triple transgenic mice expressing GCaMP6f in VIP-INs had similar expression levels and the results were combined. Average body weight at the beginning of the first session was 21.5 ± 1.3 g (PV, n = 9), 21.8 ± 4.1 g (SOM, n = 10), and 20.2 ± 2.9 g (VIP, n = 16) (mean ± SD).

METHOD DETAILS

Immunostaining and cell counting

30 μm-thick motor cortex coronal sections were prepared with a microtome (Thermo Fisher) and blocked with 10% normal goat serum, 1% bovine serum albumin and 0.3% Triton X-100 in PBS for 1 hr at room temperature. Immunostaining was then performed with 24-hr primary antibody incubation at 4°C (anti-PV (rabbit, ab11427, abcam), 1:1000; anti-GFP (chicken, GFP-1020, Aves), 1:1000, diluted in PBS with 3% normal goat serum) and 2-hr secondary antibody incubation at room temperature (anti-chicken DyLight 488 (goat, SA5-10070, Thermo Fisher), 1:1000; anti-rabbit Alexa 594 (donkey, A-21207, Thermo Fisher), 1:1000, diluted in PBS). Stained sections were mounted with CC/Mount mounting medium (Diagnostic BioSystems) and imaged with Apotome.2 (ZEISS). PV and GCaMP6f quantification was performed manually using ImageJ. Representative sections (3 for each animal) were chosen, and in each section, one rectangle area (880 μm × 680 μm) was selected for counting. We were not able to achieve reliable

labeling with SOM and VIP antibodies (T-4103 rabbit anti-Somatostatin-14 antibody (Peninsula Laboratories International, Inc.), 366004 guinea pig anti-somatostatin-28 antibody (Synaptic Systems) and 20077 rabbit anti-VIP antibody (Immunostar), all in a variety of conditions).

Surgery

Surgical procedures were performed as previously described [27]. Adult mice (6 weeks or older, male and female) were anesthetized with isoflurane and injected with Baytril (10 mg/kg), dexamethasone (2 mg/kg) and buprenorphine (0.1 mg/kg) subcutaneously to prevent infection, inflammation, and discomfort. A custom head-plate was glued and cemented to the skull. Craniotomy (~3 mm) was performed over the right caudal forelimb area (300 μm anterior and 1,500 μm lateral from the bregma). Experiments were performed at least 7 days after surgery.

Imaging

Imaging was conducted with a commercial two-photon microscope (Bscope, Thorlabs) running Scanimage using a 16x objective (NIKON) with excitation at 925 nm (Ti-Sa laser, Newport). Imaging was conducted in awake animals. For calcium imaging, images (512 × 512 pixels covering 472 × 508 μm) were recorded continuously at 28 Hz. Before each session, images at 3 different planes (10 μm step in z axis) were acquired as a reference. When ROIs were pre-defined in the previous session, it was warped with enhanced correlation coefficient maximization algorithm [28] using the reference image. ROIs were not updated during imaging. During the task, each frame was transferred to another MATLAB instance, and lateral motion was corrected by locally maximizing correlation coefficient to the reference image with gradient descent and processed for neural feedback. It was implemented with a custom MATLAB executable written in C++. Slow lateral and vertical drifts were detected by comparing correlation coefficients of the average of 200 motion corrected frames to the reference images and automatically corrected throughout imaging sessions.

Behavioral task

Before starting the neuron-pair operant conditioning task, mice first underwent 1 – 4 lick sessions when they received water drops at random intervals. These sessions allowed mice to get acclimated to the environment, and the data from calcium imaging in these sessions were used to select the target neurons. ROIs were manually selected to include the soma of neurons with at least one apparent activity event during the session. For real-time neuron-pair operant conditioning, time-varying baseline of fluorescence signal of each ROI is defined as the 8th percentile of the Savitzky–Golay filtered (order = 2, length = 9) fluorescence intensity in the preceding 200 frames. The baseline was used to calculate ΔF/F, which was then filtered by taking 15-frame moving average. Fluorescence in the background pixels (defined as all the pixels more than 2 pixels away from any ROI and also more than 32 pixels away from the border of the reference image) was averaged and the background ΔF/F was subtracted from ΔF/F of each ROI to compensate for neuropil contamination. Among candidate ROIs whose ΔF/F exceeded 0.4 at a rate between 1.5 / min to 15 / min, two were randomly selected as target ROIs (N+ and N-) for neuron-pair operant conditioning task. A threshold for the task was determined so that the simulated reward frequency using the calcium activity during lick sessions is approximately 4 / min (threshold was 0.52 ± 0.27, mean ± SD). Mice were then trained to modulate the activity of the target neurons with auditory feedback (adapted from [8]). The 15-frame rolling average of the difference of ΔF/F of N+ and N- was exponentially transformed to auditory frequency between 1 kHz and 16 kHz. 4 kHz indicates the same ΔF/F between N+ and N-, and 16 kHz indicates it reached the threshold. The delay caused by image acquisition and processing was 130 ± 33 ms (mean ± S.D., n = 40000 frames from one session). Mice were given a water reward (~10 μl) when the auditory frequency reached 16 kHz, and the sound was kept at 16 kHz for 850 ms to indicate a reward delivery (reward tone). This is followed by inter-trial intervals (ITIs) without auditory feedback, which continued until the difference of ΔF/F became smaller than 0.3 times the reward threshold for the first time after 850 ms. In neuron-pair operant conditioning sessions, imaging was performed for 6 min first without auditory feedback which served as the pre-task period. After the pre-task period, the task with auditory feedback was initiated. Behavioral sessions continued up to 50 mins or until the mice appeared disengaged from the task, e.g., not responding to available water rewards, and it was manually determined by the experimenter. During the task, the animals were monitored online by the experimenter using IR camera, and most animals showed little to no overt movements during successful trials. After 4-6 sessions of operant conditioning, a subset of mice underwent 1-3 post-operant-conditioning contingency degradation sessions. In these sessions, reward and reward tone were given at the same intervals as in the last two days of neuron-pair operant conditioning sessions without auditory feedback. In some cases the image quality degraded over the sessions, in which case the experiments were terminated and data were excluded. Only the experiments that had at least 5 imaging sessions including both neuron-pair operant conditioning and post-operant-conditioning contingency degradation sessions were included in the analysis. Some mice were trained consecutively with multiple imaging fields.

QUANTIFICATION AND STATISTICAL ANALYSES

Data analysis

Data analysis was performed using MATLAB R2014a.

Task-induced reward frequency

Instead of choosing threshold based on pre-task activity of each session as in [8], we used the same threshold over multiple sessions. Task-induced reward frequency change was calculated by taking the reward frequency between 3 and 9 mins after initiation of the

task, and subtracting the simulated reward frequency during the pre-task period. We chose this time windows because it was when mice were considered to be engaged in the task in most sessions, and the main results were not sensitive to the choice of the analysis time window (Figure S3). The simulated performance during the pre-task period did not change over time (Figure S4).

Task-induced neural activity

To delineate the strategies underlying the performance improvement in neuron-pair operant conditioning, in post hoc analysis, we computed the frequency of N+ peaks, N- peaks, N+ peaks not associated with rewards, and N- peaks during N+ inactivity. Calcium signal was considered as a peak if ΔF/F is the largest within the surrounding 1 s window and larger than the fluorescence baseline (corrected as below) by the reward threshold used for the neuron-pair operant conditioning task. Peak frequency was calculated during the time periods when the sound feedback was provided and the first 850 ms after that the sound feedback turned off after each reward. 850 ms was added to accommodate for the time that the fluorescence reaches a peak after exceeding the threshold. N+ peaks were considered associated with a reward if criteria reward was delivered within 1 s window centered at the N+ peak. N+ inactivity was when the baseline-corrected ΔF/F was below the reward threshold.

ΔF/F baseline correction

The distribution of ΔF/F tends to be skewed due to positive Ca events, and thus the mean of the ΔF/F distribution is not representative of the baseline of ΔF/F. To estimate the baseline of ΔF/F, which corresponds to the mean of ΔF/F during the inactive periods, we computed the mean of a truncated distribution defined recursively as below: First, we calculated the mean and the SD of the whole distribution during the pre-task period, and used mean ± 2 SD as the boundaries to make a truncated distribution. Next, we calculated the mean and the SD of the truncated distribution, and used mean ± 2.274 SD as the new boundaries. This step was repeated 30 times. The coefficient for SD was twice the reciprocal of the SD of Gaussian distribution ($\sigma = 1$) truncated at ± 2, which was chosen so that this method converges in one step when applied to a Gaussian distribution. For the baseline-corrected post hoc analysis described above, the estimated baseline was subtracted from the original ΔF/F.

Linear models

Task-induced frequency changes

Task-induced reward frequency changes from all the sessions were fit by a multiple linear regression model. The model formula is:

$$y \sim 1 + \text{i\_session} + (1 \mid \text{imaging field}) + (\text{i\_session-1} \mid \text{imaging field}),$$

where (1|imaging field) and (i_session-1|imaging field) indicate a random effect constant and a random effect slope term for each imaging field, i_session is a discrete variable representing the session number, imaging field is a categorical variable representing the identity of imaging fields, and y is reward or neural activity frequency adjusted to pre-task period. fitlme function of MATLAB was used to fit the model and to test the coefficient of the session number against 0.

To test the decrease of frequency of reward threshold crossing in the contingency degradation sessions, reward frequency changes from the last 2 neuron-pair operant conditioning sessions and simulated reward frequency changes from the 1-3 contingency degradation sessions were fit by a multiple linear regression model with additional task category term:

$$y \sim 1 + \text{task category} + (1 \mid \text{imaging field}),$$

where task category is a categorical variable indicating whether the session was a neuron-pair operant conditioning task or a contingency degradation session. The model was fit and the coefficient of the task category was tested against 0.

Comparison between cell types

To examine the difference between two cell types, the following model was fit to the results from the two cell types:

$$y \sim 1 + \text{i\_session} + \text{i\_session} * \text{cell type} + (1 \mid \text{imaging field}) + (\text{i\_session-1} \mid \text{imaging field})$$

When the coefficient of the interaction term is significantly different from 0, it means that the two types are significantly different.

Comparison between target types

To examine the difference due to target types (N+, N- and non-target neurons), the following model was fit to the activity of all the neurons of the two target types:

$$y \sim 1 + \text{i\_session} + \text{i\_session} * \text{target type} + (1|\text{imaging field}) + (\text{i\_session-1} \mid \text{imaging field}).$$

y indicates peak frequency (Figures 4E and 4F) and simulated reward frequency if each non-target neuron was used as N+ with the actual target neuron N- (Figure 4G), and if each non-target neuron was used as N- with the actual target neuron N+ (Figure 4H). Note that activity of individual neurons was used without averaging. When the coefficient of the interaction term is significantly different from 0, it means that the two types are significantly different.

Mediation analysis

For mediation analysis, the following three linear models were used:

$$y \sim 1 + \text{N+} + \text{N-} + (1 \mid \text{imaging field}) + (\text{N+} \mid \text{imaging field}) + (\text{N-} \mid \text{imaging field}),$$

$$N+ \sim 1 + \text{i\_session} + (1 \mid \text{imaging field}) + (\text{i\_session-1} \mid \text{imaging field}),$$

$$N- \sim 1 + \text{i\_session} + (1 \mid \text{imaging field}) + (\text{i\_session-1} \mid \text{imaging field}).$$

N+ and N- indicates N+ and N- activity frequency.

Correlation coefficient

Correlation coefficient of ΔF/F during pre-task period on the first session between neurons which met activity criteria was compared across cell types. To remove the effect of activity frequency, the following linear model was used:

$$\text{Correlation coefficient} \sim 1 + \text{mean activity frequency} + \text{cell type},$$

Where mean activity frequency is the mean of the activity frequency of the two cells of a pair, and cell type is a categorical variable representing a cell type of the pair.

Statistics

Linear models with random effect were fit to data using fitlme function in MATLAB R2014, and the significance of the coefficient is calculated by fitlme using t test. For mediation analysis, t-statistics from fitlme was used with Mediation Analysis p value calculator with normal approximation [9]. The experiments and the analysis were not blinded. Animal assignment was not randomized. Statistical sample-size estimation was not performed.

DATA AND SOFTWARE AVAILABILITY

Code for online image processing is available at https://github.com/amitani/2pNFB.