Title

Volitional modulation of optically recorded calcium signals during neuroprosthetic learning.



Journal

Nature Neuroscience, 17(6)

Authors

Clancy, Kelly

Koralek, Aaron

Costa, Rui

et al.

Publication Date

2014-06-01

DOI

10.1038/nn.3712

Peer reviewed

HHS Public Access

Author manuscript

*Nat Neurosci*. Author manuscript; available in PMC 2015 March 17.

Published in final edited form as:

*Nat Neurosci*. 2014 June ; 17(6): 807–809. doi:10.1038/nn.3712.

# VOLITIONAL MODULATION OF OPTICALLY RECORDED CALCIUM SIGNALS DURING NEUROPROSTHETIC LEARNING

**Kelly B. Clancy**[#1], **Aaron C. Koralek**[#2], **Rui M. Costa**[6], **Daniel E. Feldman**[2,3], and **Jose M. Carmena**[2,4,5]

[1]Biophysics Program, UC Berkeley, University of California, Berkeley

[2]Helen Wills Neuroscience Institute, UC Berkeley, University of California, Berkeley

[3]Dept. of Molecular & Cell Biology, UC Berkeley, University of California, Berkeley

[4]Department of Electrical Engineering and Computer Sciences, University of California, Berkeley

[5]Joint Graduate Group in Bioengineering UCB/UCSF, University of California, Berkeley

[6]Champalimaud Neuroscience Programme Champalimaud Center for the Unknown, Lisbon, Portugal

[#] These authors contributed equally to this work.

## Abstract

Brain-machine interfaces are not only promising for neurological applications, but also powerful for investigating neuronal ensemble dynamics during learning. We trained mice to operantly control an auditory cursor using spike-related calcium signals recorded with 2-photon imaging in motor and somatosensory cortex. Mice rapidly learned to modulate activity in layer 2/3 neurons, evident both across- and within-sessions. Learning was accompanied by striking modifications of firing correlations within spatially localized networks at fine scales.

Brain-machine interfaces (BMIs) have gained great momentum as a therapeutic option for patients with limb loss or immobility[1-4]. In addition, BMI tasks provide a powerful approach to study sensorimotor learning, as they enable arbitrary mapping between neuronal activity, behavioral output, and reward[5]. Recent work used BMI to demonstrate network adaptations in response to output perturbations[6], including specific functional changes in output-relevant neurons[7,8]. However, traditional BMIs based on spatially sparse electrode recordings lack fine-scale spatial information about local networks. To address this issue, we developed a BMI task in awake, head-restrained mice using 2-photon calcium imaging to record activity from every neuron in a small field of view (150 by 150 microns). We used this novel calcium-based BMI paradigm (CaBMI) to probe fine-scale network reorganization in

Corresponding authors: ruicosta@fchampalimaud.org; dfeldman@berkeley.edu; jcarmena@berkeley.edu.

AUTHOR CONTRIBUTIONS

K.B.C., A.C.K., R.M.C., D.E.F. and J.M.C. designed the experiments. K.B.C. and A.C.K. built the experimental apparatus, performed the experiments and analyzed the data. K.B.C. performed surgical procedures. A.C.K. and K.B.C. wrote the paper. K.B.C., A.C.K., R.M.C., D.E.F. and J.M.C. revised the paper.

COMPETING INTERESTS STATEMENT

The authors declare no competing financial interests.

cortical layer (L) 2/3 of both primary motor (M1) and somatosensory (S1) cortices during BMI learning.

We trained ten mice expressing the genetically-encoded calcium indicator gCaMP6f in L2/3 of M1 or S1 to modulate neural activity in response to auditory feedback (Supp. Figure 1a, Supp. Movies 1-2; Methods). This task was adapted from one used previously with electrode-based recordings[9]. Each day, two ensembles containing 1-11 neurons each were chosen to control the task (Figure 1a). The ensembles opposed each other, such that increased activity in one ensemble ("E1") above its baseline increased the pitch of the auditory feedback, while increased activity in the other ensemble ("E2") decreased the pitch. Reward was delivered when a high-pitched target was reached within 30 sec of trial initiation (hit). Incorrect trials (no target within 30 s) were signaled with white noise.

Mice learned the task rapidly (Figure 1b), with initial rapid improvement (1-3 days) followed by slower improvement (4-8 days). Mice performed above chance level after 1 day of training (Figure 1b, shaded region, N=10 mice, p=0.0036 on day 2, t(8)=4.07). Similar learning occurred using M1 or, more surprisingly, S1 (Supp. Figure 1 b-c). Hit rate increased significantly within each daily session (Supp. Figure 1d, N=72 sessions, $p=2.6\times10^{-5}$, t(43)=4.7, $R^2=0.34$). Mice reached a criterion performance level (50% hits) faster across days of training (Supp Figure 1e, N=8 days, p=0.0247, t(6)=2.98, $R^2=0.596$), suggesting that within-session learning occurs faster as between-session learning progresses. As seen previously[9], performance was not impaired by lidocaine injection into the contralateral mystacial pad (N=4 sessions, p=0.876, t(3)=0.17), and gross movements were absent preceding target hits, indicating that performance does not rely on natural movement and that neural activity, particularly in S1, is not driven by whisker reafference (Supp. Figure 2).

We next asked whether these modulations were sensitive to the action-outcome contingency[10]. After mice successfully learned the task, we ceased rewarding target hits and instead delivered rewards under a variable interval schedule (contingency degradation). Mice quickly ceased responding (Figure 1c-d; N=5 mice, p=0.0089, t(4)=4.76). When reward was reinstated using the same E1 and E2 ensembles, mice again performed at normal levels (Figure 1c; N=4 mice, p=0.791, t(3)=0.289). Thus, performance was sensitive to reward contingency. Post-hoc analysis of imaging data showed that E1 activity increased during task performance and decreased during degradation (Figure 1d). On a separate day, we performed a contingency reversal (N=3 mice) in which E1 and E2 identities were reversed from one day (day CR1) to the next (day CR2), requiring mice to reverse ensemble activity patterns to obtain reward (Supp. Figure 3a). Early during CR2, E2 in one example mouse showed clear bursting activity (consistent with its identity as E1 on CR1), and E1 showed little activity (consistent with its identity as E2 on CR1). This pattern quickly reversed as the mouse learned the new contingency (Supp. Figure 3a). We compared the hit rate on CR2 in one animal to a simulated hit rate based on the E1/E2 identity and transform algorithm from day CR1. The simulation showed initially high performance that then dropped to zero, indicating that this mouse initially performed according to the learned CR1 transform, but quickly adapted to the new CR2 transform (Figure 1e). Across all mice, the ratio of E1/E2 activity increased during CR2 (Supp. Figure 3b, n=3, R2=0.832, p=0.02,

t(18)=13.32), demonstrating that mice learn to flexibly up-modulate E1 over E2. Together, these data indicate that mice can modulate calcium signals in a contingency-dependent manner, and that these modulations can be applied to arbitrarily chosen cells.

We next investigated neural changes during learning. Mean ΔF/F increased for E1 cells over the course of individual sessions (Figure 2a; N=20 time points, $p=1.17\times10^{-11}$, t(18)=15.09, $R^2=0.927$), decreased during subsequent contingency degradation (Supp Figure 4a; N=20 time points, p=0.05, t(18)=2.08), and increased again during reinstatement, mirroring changes in hit rate (Supp. Figure 4b, c). In contrast, mean ΔF/F did not significantly change for E2 cells (Supp. Figure 4d, e). This may reflect a bias toward volitional increases, rather than decreases, of L2/3 activity.

Calcium imaging detects activity even in neurons that are rarely active, which are numerous in L2/3[11,12]. These cells are undersampled by extracellular recordings and are often neglected in BMI studies[13]. There was a 30-fold range of baseline spontaneous activity across L2/3 cells (Fig. 2b). We found the most dramatic increases in task-related activity in E1 cells with initially low baseline. E2 cells tended to lower the frequency of calcium events during task engagement, independent of baseline activity (Figure 2b), though mean activity remained unchanged (Supplementary Figure 4d). Thus, task learning preferentially recruited low-active E1 neurons to become more active. These neurons clearly contributed to learning, because learning occurred normally when all E1 cells had low or zero baseline activity (Supplementary Figure 5a, N=46 sessions, p=0.83, t(44)=1.56), suggesting a role for "silent" L2/3 neurons in learning[12]. We also found that within multi-cell E1 ensembles, multiple cells increased fluorescence around hits, including low- and zero-baseline cells, and so performance was not carried by single neurons (Supplemental Figure 5b).

To examine higher-level network dynamics during learning, we first calculated mean cross-correlation histograms time-locked to the occurrence of large fluorescence events in either E1 or E2 ("output cells"; Methods). Output cells developed coordinated, synchronous activity with other cells in the same ensemble (Figure 2c,d). E2 also developed a tendency to spike before E1 (Figure 2c, d), likely reflecting a strategy of bursting E2 for trial initiation, followed by bursting of E1 for target achievement. This coordinated activity was not present in non-output cells that were simultaneously imaged (Figure 2c,d; "indirect cells"). This prompted us to investigate correlations between cells over the course of individual sessions. Correlations between output cells in the same ensemble increased significantly during the session (N=5 time points, p=0.0198, t(3)=4.55, $R^2=0.874$), while correlations between indirect cells did not (Figure 2e), and this enhancement was observable in individual animals (Supplementary Figure 6a). Output cells also became more correlated over days of training, even though neural composition of ensembles changed (Supplementary Figure 6b). This is analogous to increased correlations of functionally-related cells during motor learning[14] and could reflect millisecond-precision coupling that has been demonstrated with electrodes[15].

We next examined how fine-scale (~10-100 microns) spatial organization of ensembles impacted learning. Performance did not vary systematically with distance between output ensembles (measured by E1 and E2 centroids) (N=71 units, p=0.955, t(69)=0.056), but did vary with size of ensembles: animals performed better with fewer neurons, suggesting that it

was difficult to maintain coordinated control over large ensembles (Supplemental Figure 7). Additionally, high baseline correlations between ensembles predicted worse performance (Supplementary Figure 8).

Learning was accompanied by interesting dynamics within local networks surrounding the output ensembles. For each indirect cell, we calculated the correlation between its mean fluorescence and a moving average of the animal's instantaneous hit rate. We found that activity in indirect cells near E1 (<50 μm away) was significantly more correlated with hits than activity in distant indirect cells (>100 μm away; Figure 2f, N=251 cells, p=0.048, t(249)=1.98). Finally, we calculated mean target-related modulations in indirect cells for early and late epochs within daily sessions. Early in sessions, indirect cells near E1 exhibited increased ΔF/F around hits, while distant indirect cells did not. In contrast, indirect cells near E2 showed modest suppression early in sessions (Figure 2g). Thus, early in the session mice modulate activity in a relatively large local network surrounding output cells[7], but as the session progresses this target-related modulation in indirect cells disappears, such that mainly output cells exhibited task-related modulations. This suggests that mice are able to hone in on individual output cells during learning and precisely modulate these cells for efficient target achievement. Indirect neurons that were more highly spontaneously correlated with E1 cells, and therefore more likely to be embedded in the same local sub-network[16], exhibited significantly less dampening in task-related modulations (Supplementary Figure 9a, b; N=851 cells, p = 9e-9, p = 0.38, respectively). Given the rapid falloff of spontaneous correlations with distance, such fine-scale effects might be undetectable by electrode-based methods (Supplementary Figure 9c). This spatial restriction in activity is similar to sparsening of cortical representations during classical conditioning[17].

To our knowledge, this CaBMI study is the first demonstration that mice can volitionally modulate calcium dynamics in L2/3 of M1 and S1, and the use of imaging enabled dissection of learning-related network modifications during BMI with unprecedented spatial resolution (~10-100 μm). Importantly, this novel paradigm provides a powerful tool for investigating the spatial extent of functional and structural plasticity during neuroprosthetic learning.

## METHODS

All animal procedures were performed in accordance with UC Berkeley Animal Care and Use Committee regulations. 6 C57BL/6J and 4 CD1 male wild-type mice were used in these experiments, ranging in age from postnatal day 30-45. Animals were housed with a 12h dark – 12h light reversed light cycle. All behavioral tests were performed in the same cohort of animals.

### Surgery

Mice were anesthetized using isoflurane (2% vol isoflurane/vol O2) and placed in a stereotaxic apparatus. Body temperature was maintained at 37°C using a feedback-controlled heating pad (FHC, 40-90-8D) and a small incision was made in the scalp. The skull was cleaned and a steel headplate was affixed over M1 (1 mm rostral, 1 mm lateral to Bregma) or S1 (1 mm caudal, 3 mm lateral to Bregma) using Metabond dental cement

(Parkell, S380). A 3 mm craniotomy was opened over M1 or S1, and 200 nL of AAV2.9 Syn.GCamp6f.WPRE.SV40[20] (University of Pennsylvania Vector Core) was injected 250 μm below the pia using a Nanoliter 2000 injector (World Precision Instruments). The tracer was delivered using a pulled glass pipette (tip diameter 40-60 μm) at a rate of 40 nL/minute. The pipette was left in the brain for 10 minutes after completion of the injection to prevent backflow. After the pipette was removed, the brain was covered with silicone oil (Sigma product # 181138) and a glass coverslip was affixed to the skull with dental cement, as previously described[21]. We allowed 2 weeks for recovery and gCaMP6f expression.

## Two-Photon Imaging

*In vivo* imaging was performed with a Moveable Objective Microscope (Sutter) using a Chameleon Ultra Ti:Sapphire mode-locked laser (Coherent, Santa Clara CA) tuned to 900 nm. Photons were collected with a Hamamatsu photomultiplier tube (H10770PA-40) using a Nikon objective (16x, 0.8 NA). Animals were head-fixed on a custom-made spring mounted imaging platform and placed under the 2p microscope. This setup allowed them to run freely, and their movements were recorded by an accelerometer fixed to the underside of the platform. Frames of 128×512 pixels (~160 × 160 μm) were collected at 7.23 Hz using ScanImage software[22] at 130-180 μm below the pia. The same imaging fields were used every day, localized by landmarks in the surface blood vessels. Imaged fields were stable over the course of training, and because the cortex was stabilized by a snugly fitting coverslip, only severe movements caused motion artifacts. Motion correction for slow drift in the imaging field was performed manually. Any period of gross movement during the task that caused cells to move out of their ROIs resulted in poor task performance, as ΔF/F of E1 was reduced. In this sense, mice were punished for excessive movement and seem to have learned to remain still during the task (Supplementary Figure 2c).

## Behavioral Task

Two ensembles of 1-11 single cells each were chosen for inclusion in the “output” population. Cells with bright nuclei, indicating over-expression, were excluded, as were cells with many, poorly separable calcium events, an activity pattern indicative of fast-spiking interneurons. No other selection criteria were used to partition the recorded cells into each ensemble. We also ensured that many cells with good signal were included in the indirect population to enable a proper comparison. The cells assigned to the output population were changed on some days.

Ensemble activity was measured as mean ΔF/F for all component neurons. Fluorescence values from these ensembles were binned in 200 millisecond bins and entered into an online transform algorithm that related neural activity to the pitch of an auditory cursor. By modulating activity in these ensembles, rodents controlled the pitch of the cursor. The modulations that we required of the mice were calibrated daily based on a baseline recording session of roughly 2 minutes. Next, 10-15 minutes of spontaneous baseline activity was recorded to assess chance levels of performance and spontaneous levels of activity. Fluorescence values were smoothed by a moving average of the past 3 time points. Changes in the frequency of the auditory cursor were binned in quarter-octave intervals to match rodent psychophysical discrimination thresholds[23]. Mice then had to modulate calcium

dynamics in these neuronal ensembles to move the cursor to a high-pitched target tone that was associated with a 10% sucrose solution reward. A trial was marked incorrect if a target was not achieved within 30 seconds of trial initiation. A trial was self-initiated when E1 and E2 activity returned to baseline levels (either by decreased activity in E1 or increased activity in E2), which reset the tone to its starting pitch.

Regions of interest (ROIs) were extracted from recorded neural data in real time. These ROIs were entered into custom routines in MATLAB (Mathworks, Natick, MA) that translated fluorescence levels into the appropriate feedback pitch and played the pitch on speakers mounted on 2 sides of the imaging platform. Frequencies used for auditory feedback ranged from 1-24 kHz in quarter-octave increments. When a target was hit, a MATLAB-controlled Data Acquisition board (National Instruments, Austin, TX) triggered the operant box to supply the appropriate reward to rodents. Each daily training session lasted 48±2 min (71±4 trials).

## Data Analysis

All analyses were performed with custom written routines in MATLAB (Mathworks, Natick, MA). Recorded movies were spatially aligned using the dftregistration routine in MATLAB[24]. Regions of interest were manually selected to include the soma of neurons that appeared consistently throughout all recorded movies. Fluorescence traces were extracted from each ROI and data is presented as the relative change in fluorescence, $\Delta F/F$.

No statistical methods were used to pre-determine sample sizes but our sample sizes are similar to those generally employed in the field.

For analyses of behavioral performance during the contingency degradation and reinstatement, the first 10 trials of a session were removed before calculating performance.

For all sliding window analyses, sessions were divided into an equal number of bins to determine the window size, and the step size was a fraction of this window size.

For the cross-correlation histograms, fluorescence traces from output cells were z-scored and values above 3 standard deviations were considered an event. The first time point in which fluorescence values crossed this threshold during each event was used for time-locking. Fluorescence values in other populations of cells were then averaged around these indices.

In all cases, multiple comparisons were controlled for using the Bonferroni correction. Differences between groups were tested with T-tests and trends over time were tested with T-tests on the linear regression coefficients. All statistical tests were two-tailed.

Data distributions were assumed to be normal, but this was not formally tested. Data collection and analysis were not performed blind to the experimental conditions. Randomization was not performed, as the experiment primarily involved within-animal comparisons and there were not multiple experimental cohorts.

## Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

## ACKNOWLEGEMENTS

This work was supported by the NIH grant 1R01NS072416-01 to D.E.F., the NSF grant CBET-0954243 to J.M.C., the European Research Council grant 243393 to R.M.C., and the NSF Graduate Research Fellowship grant DGE 1106400 to K.B.C. We thank V. Jayaraman, R. Kerr, D. Kim, L. Looger, and K. Svoboda for generously sharing gCaMP material, and V. Athalye and K.D. Smith for technical assistance.

## REFERENCES

1. Nicolelis MA. L. Nature. 2001; 409:403–407. [PubMed: 11201755]
2. Carmena JM, et al. PLoS Biology. 2003; 1:e2. [PubMed: 14624234]
3. Hochberg LR, et al. Nature. 2012; 485:372–375. [PubMed: 22596161]
4. Collinger JL, et al. The Lancet. 2013; 381:557–564.
5. Green AM, Kalaska JF. Trends in Neurosciences. 2011; 34:61–75. [PubMed: 21176975]
6. Jarosiewicz B, Schummers J, Malik WQ, Brown EN, Sur M. Current Biology. 2012; 22:269–277. [PubMed: 22305753]
7. Ganguly K, Dimitrov DF, Wallis JD, Carmena JM. Nature Neuroscience. 2011; 14:662–667.
8. Koralek AC, Costa RM, Carmena JM. Neuron. 2013; 79:865–872. [PubMed: 23954030]
9. Koralek AC, Jin X, Long JD II, Costa RM, Carmena JM. Nature. 2012; 483:331–335. [PubMed: 22388818]
10. Hilario MRF, Clouse E, Yin HH, Costa RM. Frontiers in Integrative Neuroscience. 2007; 1:1–12. [PubMed: 18958229]
11. O'Connor DH, Peron SP, Huber D, Svoboda K. Neuron. 2010; 67:1048–1061. [PubMed: 20869600]
12. Barth AL, Poulet JFA. Trends in Neurosciences. 2012; 35:345–355. [PubMed: 22579264]
13. Shoham S, O'Connor DH, Segev R. J Comp Physiol A. 2006; 192:777–784.
14. Komiyama T, et al. Nature. 2010; 464:1182–1186. [PubMed: 20376005]
15. Engelhard B, Ozeri N, Israel Z, Bergman H, Vaadia E. Neuron. 2013; 77:361–375. [PubMed: 23352171]
16. Harris KD, Mrsic-Flogel TD. Nature. 2013; 503:51–58. [PubMed: 24201278]
17. Gdalyahu A, et al. Neuron. 2012; 75:121–132. [PubMed: 22794266]
18. Dias-Ferreira E, et al. Science. 2009; 325:621–625. [PubMed: 19644122]
19. Costa RM. Current Opinion in Neurobiology. 2011; 21:579–586. [PubMed: 21641793]
20. Chen T-W, et al. Nature. 2013; 499:295–300. [PubMed: 23868258]
21. Holtmaat A, et al. Cold Spring Harbor Protocols. 2012; 2012:694–701. [PubMed: 22661440]
22. Pologruto TA, Sabatini BL, Svoboda K. BioMed Eng OnLine. 2003; 2:13. [PubMed: 12801419]
23. Han YK, Köver H, Insanally MN, Semerdjian JH, Bao S. Nature Neuroscience. 2007; 10:1191–1197.
24. Guizar-Sicairos M, Thurman ST, Fienup JR. Optics Letters. 2008; 33:156–158. [PubMed: 18197224]

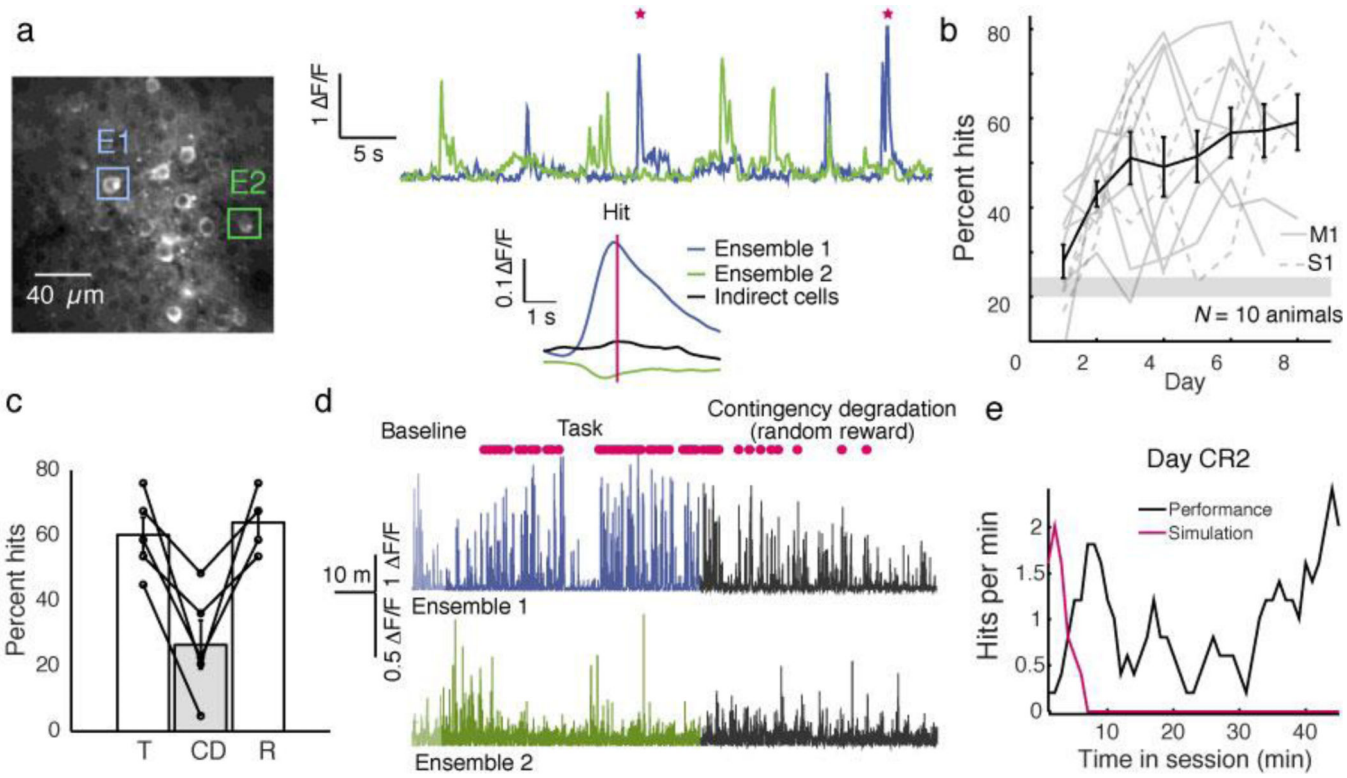


**Figure 1. Mice learn to intentionally modulate calcium dynamics**

**a.** Example imaging field (left) and recordings from cells in E1 and E2 (top right). Red stars indicate hits. Bottom right, mean ensemble fluorescence around hits. **b.** Performance over 8 days of training. Mean performance is shown in black, individual animals in gray. Error bars denote s.e.m. Shaded region denotes chance performance. **c.** Performance rapidly dropped compared to normal task levels (T) during the contingency degradation (CD). Performance returned to previous levels during reinstatement (R). **d.** E1 ΔF/F increases during the task and decreases during CD. Likewise, target hits (red) increase in frequency over training, and decrease during CD. **e.** At the beginning of day CR2, the animal initially performs as if the previous day's transform algorithm were still in use, but quickly learns the new transform.

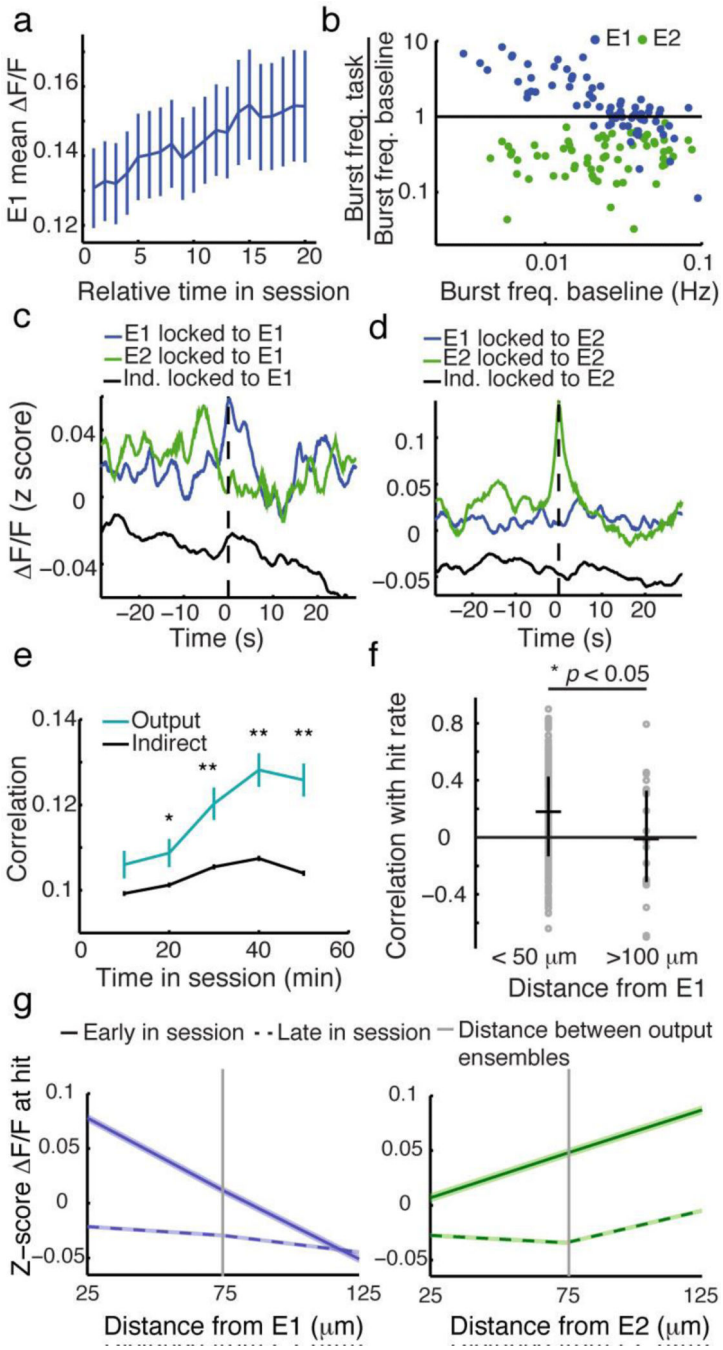


**Figure 2. Local network reorganization accompanies neuroprosthetic learning**

**a.** Mean fluorescence increases in E1 cells over the course of a session. Inset: Mean target-locked fluorescence. **b.** E1 cells with low baseline activity increase their activity more during the task than cells with high baseline activity. E2 cells suppress their activity evenly. Note logarithmic scale. **c.** Activity in E1, E2 and indirect cells time-locked to large events in E1 cells. **d.** Activity in E1, E2 and indirect cells time-locked to large events in E2 cells. **e.** Correlations increase between output cells (cyan) during the session, with no similar increase in correlations between indirect cells (black). **f.** Indirect cells near output cells have

more task-related activity than those far from output cells. Circles are individual cells, bars indicate s.e.m. **g.** Early in a session (solid lines), target-related modulations in indirect cells decrease with distance from E1 cells (blue) and increase with distance from E2 cells (green). Later in the session (dashed lines) there are no significant modulations in indirect cells, regardless of distance from output cells.