



New neural activity patterns emerge with long-term learning

Emily R. Oby^{a,b,c,d,e}, Matthew D. Golub^{b,f,g,h}, Jay A. Hennig^{b,i,j}, Alan D. Degenhart^{a,b,c,d}, Elizabeth C. Tyler-Kabara^{a,k,l,m}, Byron M. Yu^{b,f,i,n,1}, Steven M. Chase^{b,i,n,1}, and Aaron P. Batista^{a,b,c,d,1,2}

^aDepartment of Bioengineering, University of Pittsburgh, Pittsburgh, PA 15213; ^bCenter for the Neural Basis of Cognition, University of Pittsburgh and Carnegie Mellon University, Pittsburgh, PA 15213; ^cUniversity of Pittsburgh Brain Institute, Pittsburgh, PA 15213; ^dSystems Neuroscience Center, University of Pittsburgh, Pittsburgh, PA 15213; ^eDepartment of Neurobiology, University of Pittsburgh School of Medicine, Pittsburgh, PA 15213; ^fDepartment of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213; ^gDepartment of Electrical Engineering, Stanford University, Stanford, CA 94305; ^hWu Tsai Neurosciences Institute, Stanford University, Stanford, CA 94305; ⁱCarnegie Mellon Neuroscience Institute, Carnegie Mellon University, Pittsburgh, PA 15213; ^jMachine Learning Department, Carnegie Mellon University, Pittsburgh, PA 15213; ^kDepartment of Neurological Surgery, University of Pittsburgh, Pittsburgh, PA 15213; ^lDepartment of Physical Medicine and Rehabilitation, University of Pittsburgh, Pittsburgh, PA 15213; ^mMcGowan Institute for Regenerative Medicine, University of Pittsburgh, Pittsburgh, PA 15213; and ⁿDepartment of Biomedical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213

Edited by Michael E. Goldberg, Columbia University College of Physicians, New York, NY, and approved May 16, 2019 (received for review December 5, 2018)

Learning has been associated with changes in the brain at every level of organization. However, it remains difficult to establish a causal link between specific changes in the brain and new behavioral abilities. We establish that new neural activity patterns emerge with learning. We demonstrate that these new neural activity patterns cause the new behavior. Thus, the formation of new patterns of neural population activity can underlie the learning of new skills.

skill learning | neural population | motor cortex | brain–computer interface

Our understanding of learning is grounded in the concepts of synaptic plasticity (1, 2) and cortical map plasticity (3, 4). However, we lack an explanation for how such changes give rise to new behavioral capacities. Establishing a causal link from learning-related changes in the brain to new behavioral capacities would require knowing which neurons drive behavior, as well as the relationship between the activity of those neurons and behavior. Then, any observed change in behavior can be attributed to an observed change in the neural activity. A brain–computer interface (BCI) enables us to link changes in neural activity directly to learning because the relationship between the neural activity and the behavior is known exactly, and only the neurons we record directly influence behavior (5).

We hypothesize that learning new skills can be achieved by the formation of new patterns of neural activity, where patterns of neural activity are defined as the joint firing rate of a population of neurons measured during a brief time window (*SI Appendix, Materials and Methods*). To demonstrate that new neural activity patterns drive learning, we must achieve three objectives: First, we need to encourage new neural activity patterns to form. Second, we need to detect new patterns, should they appear. Third, we need to show that the new patterns directly cause the new behavioral capacities that emerge with learning. A BCI learning paradigm provides a framework whereby we can achieve these objectives.

In our BCI paradigm, a monkey moves a computer cursor from the center of a screen to one of eight peripheral targets by volitionally modulating the activity of a population of ~90 neural units recorded with a multielectrode array chronically implanted in the arm region of primary motor cortex (Fig. 1*A* and *SI Appendix, Materials and Methods*). Each experiment begins by presenting the monkey with an “intuitive mapping” that relates neural activity patterns to cursor velocities in a way that provides proficient control without requiring the animal to learn. We then induce learning by presenting the monkey with a novel mapping from neural activity to cursor velocity (6–8).

Previously, we have shown that the structure of neural population activity limits the learning that can occur within a single day (8). We characterize how neurons naturally covary as the “intrinsic manifold” (Fig. 1*B*). In our BCI system, neural

activity is mapped into the intrinsic manifold and then to cursor velocity (Fig. 1*A* and *B*). This enables us to study learning by constructing two types of novel BCI mappings. A novel BCI mapping that is consistent with the intrinsic manifold, i.e., a “within-manifold perturbation” (WMP), can be well learned within a single day. The neural strategy for learning WMPs does not involve the formation of new neural activity patterns. Instead, learning occurs by reassociating preexisting patterns of neural activity with different intended movements (9). In contrast, novel BCI mappings that are inconsistent with the intrinsic manifold, i.e., “outside-manifold perturbations” (OMPs), are not well learned within a single day. OMPs encourage the formation of new neural activity patterns. Thus, we achieve the first objective of our study by challenging animals to learn to use OMP mappings given extended practice over several days. To this end, we use the framework of the intrinsic manifold to repeatedly and reliably construct novel BCI mappings that encourage the formation of new patterns of neural activity and ask whether the new mappings are learnable. Through subsequent analyses, we achieve the second and third objectives of our study, demonstrating that OMP mappings

Significance

Consider a skill you would like to learn, like playing the piano. How do you progress from “Chopsticks” to Chopin? As you learn to do something new with your hands, does the brain also do something new? We found that monkeys learned new skilled behavior by generating new neural activity patterns. We used a brain–computer interface (BCI), which directly links neural activity to movement of a computer cursor, to encourage animals to generate new neural activity patterns. Over several days, the animals began to exhibit new patterns of neural activity that enabled them to control the BCI cursor. This suggests that learning to play the piano and other skills might also involve the generation of new neural activity patterns.

Author contributions: E.R.O., M.D.G., A.D.D., B.M.Y., S.M.C., and A.P.B. designed the experiments; E.R.O. performed the experiments; E.R.O., M.D.G., J.A.H., A.D.D., B.M.Y., S.M.C., and A.P.B. designed the analyses and discussed the results; E.R.O. analyzed the data; E.R.O. wrote the paper; E.R.O. and E.C.T.-K. performed animal surgeries; and all authors commented on the manuscript.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹B.M.Y., S.M.C., and A.P.B. contributed equally to this work.

²To whom correspondence may be addressed. Email: aaron.batista@pitt.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1820296116/-DCSupplemental.

Published online June 10, 2019.

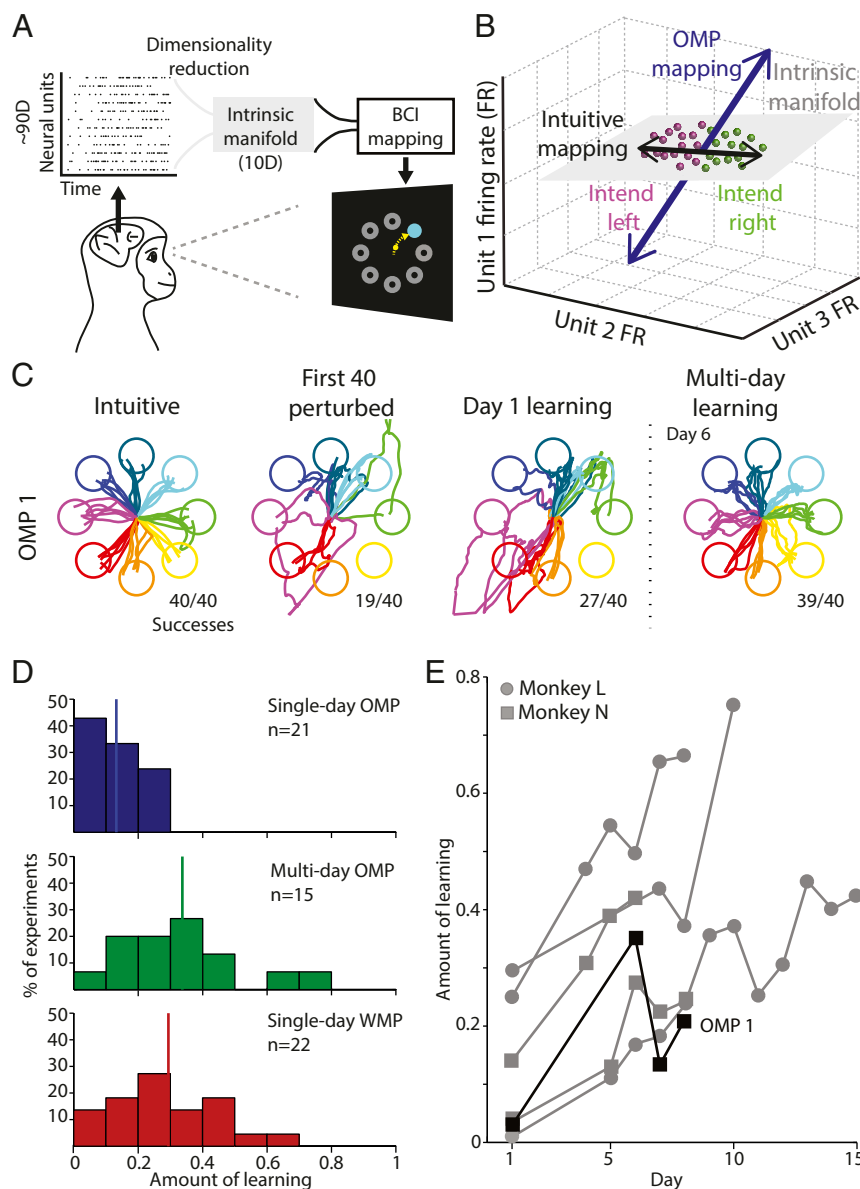


Fig. 1. Using a brain–computer interface to study learning. (A) Schematic of the BCI system. Monkeys controlled a BCI cursor (yellow) to acquire one (cyan) of eight possible (gray) targets by modulating their neural activity. In our BCI system, $\sim 90\text{D}$ neural activity is first mapped into the 10D intrinsic manifold, and then to 2D cursor velocity. (B) A simplified, conceptual schematic of neural activity patterns (dots). The neural activity patterns tend to lie in a low-dimensional subspace, termed the intrinsic manifold (gray plane). Monkeys move the BCI cursor by volitionally modulating their neural activity. At the beginning of each experiment, cursor velocities were determined by an intuitive BCI mapping (black arrow). Here, patterns colored green would move the cursor to the right and purple patterns would move the cursor to the left. To induce learning, we changed the mapping to an OMP mapping (blue arrow). Under the perturbed mapping, neural activity patterns map to different cursor velocities than under the intuitive mapping. This encourages the monkeys to learn. (C) Cursor trajectories for successful trials during a representative multiday OMP experiment (i.e., OMP 1; the multiday OMP experiment beginning on June 17, 2016). “Intuitive” trajectories show 40 consecutive trials with the intuitive mapping as an example of proficient cursor control. The second column shows the first 40 trials after switching to the OMP mapping. Performance is impaired. The third column shows the best 40 consecutive trials on day 1. The fourth column shows the best 40 consecutive trials after 6 d of practice. (D) Quantifying the amount of learning for single-day OMP (blue), multiday OMP, incremental training (green), and single-day WMP (red) experiments. An amount of learning of 1 indicates complete learning, and a value of 0 indicates no learning. Vertical lines indicate the mean of each distribution. (E) Learning curves for the six experiments with the greatest amount of learning. The example in C is highlighted in black.

are indeed learned by the formation of new patterns of neural activity, and that those patterns directly drive the new behavior.

Results

We conducted 15 multiday OMP learning experiments (ranging from 6 to 16 d per experiment, average 9.2 d) across two monkeys. Each of these novel OMP mappings was learned over several days (Fig. 1 C and D). With multiday practice, the amount of OMP learning is substantially greater than single-day

OMP learning (Fig. 1D; t test $P < 10^{-4}$), and is comparable to the single-day learning we previously observed for WMPs (8) (Fig. 1D; t test $P = 0.53$). To facilitate learning, we employed an incremental training paradigm (10) (*SI Appendix, Materials and Methods* and Fig. S1 A–D). Multiday exposure to an OMP with no incremental training led to inconsistent learning (*SI Appendix, Fig. S1 E–G*). The incremental training approach was not effective within a single day (*SI Appendix, Fig. S1G*). For most experiments, learning proceeded in a manner that resembles skill

learning (11): gradual improvement over the course of many days, with some dips and rebounds in performance (Fig. 1E and *SI Appendix*, Fig. S1 C and D). The dips and rebounds likely reflect some combination of the natural skill-learning process, motivation during a difficult task, and day-to-day recording instabilities (12–14).

The second objective of our study is to detect whether new neural activity patterns emerged during multiday OMP learning. By construction, forming new patterns of neural activity is the optimal neural strategy for learning to control the cursor under an OMP mapping because this would lead to the fastest cursor speeds. However, it is possible that the brain is unable to form new patterns because constraints exist on the patterns of neural activity that a population of neurons can exhibit (8, 9, 15–18). If this is the case, the monkey could still show some limited behavioral improvements by learning to reassociate preexisting patterns of neural activity with different intended movements (9). Thus, behavioral improvements alone are not sufficient to conclude that new patterns of neural activity have emerged.

Detecting the appearance of a new pattern of neural activity in the high-dimensional neural space is difficult because we observe only a limited number of patterns relative to the dimensionality of the space. Instead, we leveraged our BCI framework to look

for the emergence of new neural activity patterns within the low-dimensional space specified by the BCI mapping. We started by defining the patterns of neural activity observed before learning as the “intuitive neural repertoire” (9). Then we projected the ~90D neural activity patterns comprising the intuitive neural repertoire into the 2D velocity space defined by the OMP mapping (Fig. 2A). This defines the limits on cursor velocities the monkey could produce through the OMP mapping if he only expressed patterns of neural activity from within his intuitive neural repertoire (Fig. 2B and *SI Appendix, Materials and Methods*). We term this the neural “speed limit.” Any cursor velocities we observe after learning that exceed the speed limit must have arisen from neural activity patterns that are outside of the intuitive neural repertoire, and thus are new (Fig. 2B and C). Over the course of many days, monkeys learned to move the cursor at velocities that exceeded the speed limit for many targets (Fig. 2D–F). The percentage of neural activity patterns that are new significantly increased over days (t test, $P = 0.0015$; Fig. 2G). This shows that the brain can generate new neural activity patterns when learning to perform a new skill, but that it takes several days to do so.

The third objective of our study is to show that the new neural activity patterns caused the behavioral improvements. In conventional learning studies, neural changes are observed that

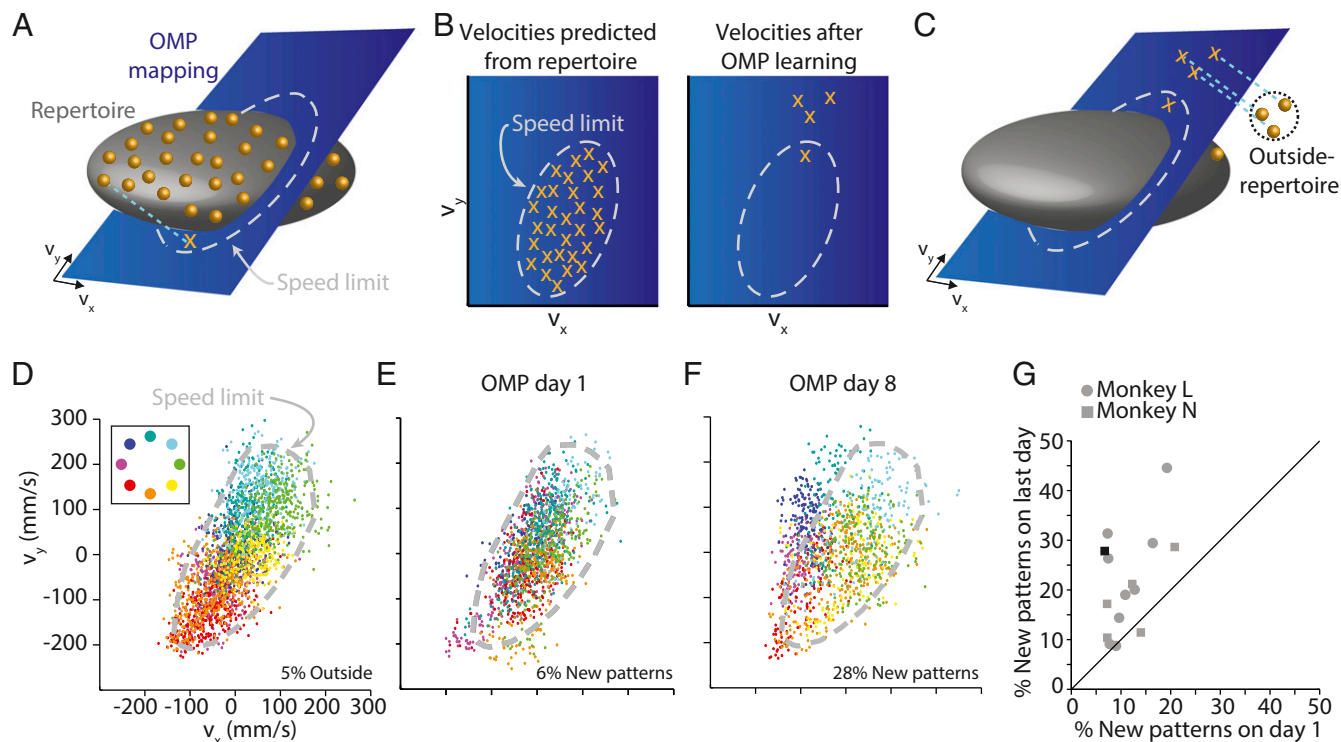


Fig. 2. New neural activity patterns emerge with long-term BCI learning. (A–C) A schematic of the technique used to identify outside-repertoire activity patterns. The OMP maps ~90D neural population activity patterns to 2D cursor velocities. Here we illustrate using 3D neural activity patterns and a 2D OMP mapping. (A) The neural activity patterns (orange dots) generated by the animal while using the intuitive mapping define an intuitive neural repertoire (dark gray ellipsoid). Each neural activity pattern maps to a cursor velocity (orange X, as one example) through the OMP mapping (blue plane). (B, Left) The velocities predicted from the intuitive neural repertoire (orange Xs) through the OMP mapping define a speed limit (dashed gray ellipse). (Right) After learning, cursor velocities are observed that exceed the speed limit (three Xs outside of gray ellipse). (C) If the monkey produces cursor velocities that exceed the speed limit, those velocities were generated by neural activity patterns that lie outside of the intuitive neural repertoire and thus are new. (D–F) Using the speed limit to detect new neural activity patterns for an example experiment (OMP 1 from Fig. 1). (D) Velocities generated from the intuitive neural repertoire mapped through the OMP mapping. Each dot is the velocity resulting from one neural activity pattern (45-ms bin). Dots are colored by instructed target location (*Inset*). The speed limit is defined as the 95% convex hull (gray dashed line). By definition, 5% of the neural activity patterns are outside of the speed limit. (E) Day 1 velocities, generated while using the OMP mapping, mostly fell within the speed limit. Dots shown are from the 40 consecutive trials when behavior was the best on day 1. On these trials, 6% of the neural activity patterns were outside of the speed limit. (F) On day 8, some velocities exceeded the speed limit (e.g., patterns corresponding to purple, blue, and teal targets). Same conventions as in E. (G) The percentage of neural activity patterns that were new on the last day of OMP learning exceeds the percentage seen on the first day of learning for most experiments. Each symbol is one multiday OMP learning experiment. OMP1 is indicated in black.

accompany learning, but it has been difficult to know if those changes are directly responsible for the learned behavior. A BCI allows us to assess the behavioral consequence of any given neural activity pattern. To assess the behavioral impact of each neural activity pattern, we measured the component of cursor velocity in the direction of the target, which we term “progress” (Fig. 3A). Higher progress indicates straighter and/or faster cursor movements. Over the course of a multiday experiment, progress improved (Fig. 3B). Increases in progress were positively correlated with the emergence of more new neural activity patterns (Fig. 3C). This indicates that the monkey learned to move the cursor faster and straighter to the target in part by producing new neural activity patterns.

Before we can conclude that new neural activity patterns emerge due to learning, we need to ensure that they do not emerge by chance or due to neural recording instabilities. We performed three key controls to address this. First, we show that the new activity patterns were specific to the OMP being learned, and not generally helpful through other OMP mappings. To test this, we mapped the neural activity patterns generated by the monkey after learning through 500 random OMP mappings to assess how much apparent learning would have occurred with mappings which the monkey never experienced. The apparent learning rarely exceeded the learning observed with the OMP mappings which the monkey did experience (*SI Appendix, Fig. S2*). This rules out the possibility that the animal learned a general strategy that would work for any OMP.

Second, we show that new patterns do not emerge without substantial learning pressure. That is, they do not appear by chance or because of neural recording instabilities. To assess this, we conducted control experiments in which the monkey used the same intuitive mapping for several days. This is a low learning pressure scenario because while control is good from the beginning of the experiment, there may still be incentive for the animal to increase his reward rate. The recording instabilities during these control experiments were indistinguishable from our multiday OMP learning experiments (*SI Appendix, Fig. S3*), and, importantly, they do not result in the same emergence of new patterns of neural activity (*SI Appendix, Fig. S4B*). We also

mapped the neural activity recorded during multiday intuitive mapping experiments through 500 random OMP mappings, which the monkey never experienced (*SI Appendix, Fig. S4C*). This is a “no learning pressure” scenario because the monkey never received feedback about how his neural activity would have impacted behavior through these mappings. The day-to-day recording instabilities in a no learning pressure scenario did not manifest as learning (*SI Appendix, Fig. S4D*).

Third, we assessed the extent to which the new activity patterns moved outside of the speed limit during learning. If new neural activity patterns are formed by expanding the neural repertoire, the distance of the new patterns from the speed limit should increase with learning. We found that in multiday OMP learning, not only does the percentage of patterns that are new increase with learning (Fig. 2 and *SI Appendix, Fig. S4A and D*), but also the distance of the new patterns from the speed limit increases (*SI Appendix, Fig. S5*). Taken together these controls confirm that the new neural activity patterns that emerge during multiday OMP learning are specific to the learned mapping, are directly responsible for behavioral improvements, and cannot be attributed to recording instabilities or other chance events.

We have shown that learning can proceed by the formation of new neural activity patterns that directly drive behavioral improvements. We wondered what characteristics make these new neural activity patterns different from the existing patterns. It might be that the new patterns are well explained by the pre-existing correlation structure, which is captured by the intrinsic manifold. Such new patterns could arise from organized changes in neural firing rates that conform to the preexisting correlation structure. We characterize this as an inside-manifold change (Fig. 4A). Alternatively, it might be that the new patterns arise from changes in the correlations between the neurons. We characterize this as an outside-manifold change (Fig. 4B). To determine the extent to which the new patterns generated by the animal resulted from firing rate changes or correlation changes, we decomposed each neural activity pattern into an inside-manifold component and an outside-manifold component (*SI Appendix, Fig. S6A and B*). This enabled us to separately determine the inside-manifold and outside-manifold contributions

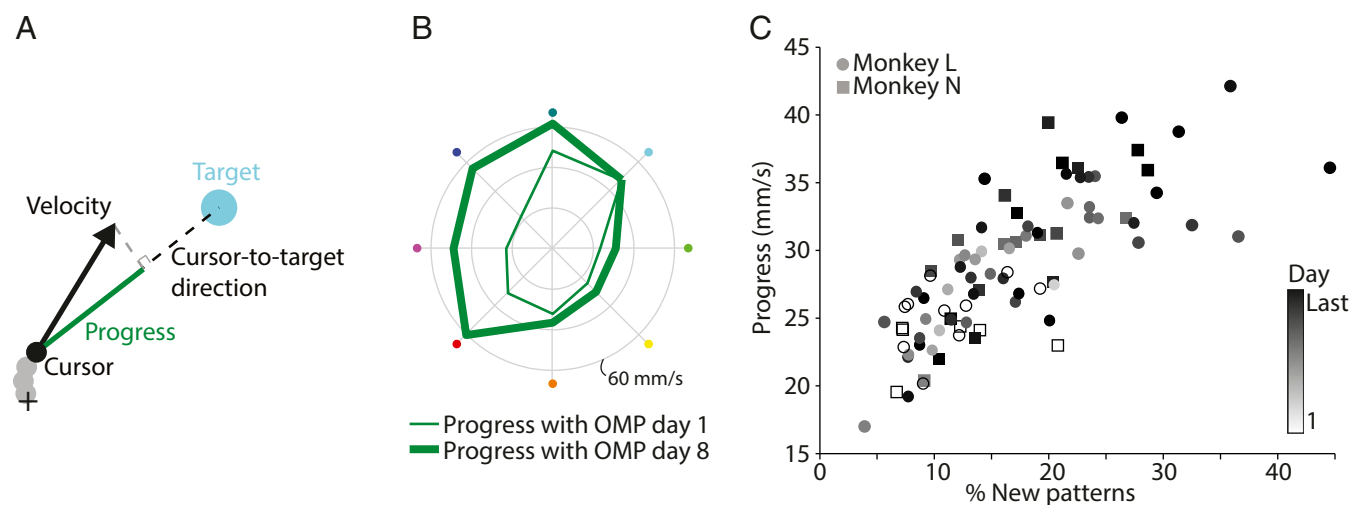


Fig. 3. New neural activity patterns drive behavioral improvements. (A) Illustration of the progress metric. Progress is defined as the component of cursor velocity in the cursor-to-target direction. The + represents the center of the screen. Gray circles are cursor positions at previous time points. (B) Mean progress toward each target on day 1 (thin) and day 8 (thick) for an example experiment (OMP 1). (C) The cursor movements showed more progress when there was a larger percentage of neural activity patterns that were new (Pearson correlation coefficient $r = 0.76$, $P = 6 \times 10^{-16}$), and generally increased with several days of practice. Each symbol is the mean progress averaged over all eight targets on 1 d of one multiday OMP experiment. The shading of the symbols indicates the day within a given multiday experiment. In general, earlier days showed relatively few new patterns, and later days showed more new patterns and better progress.

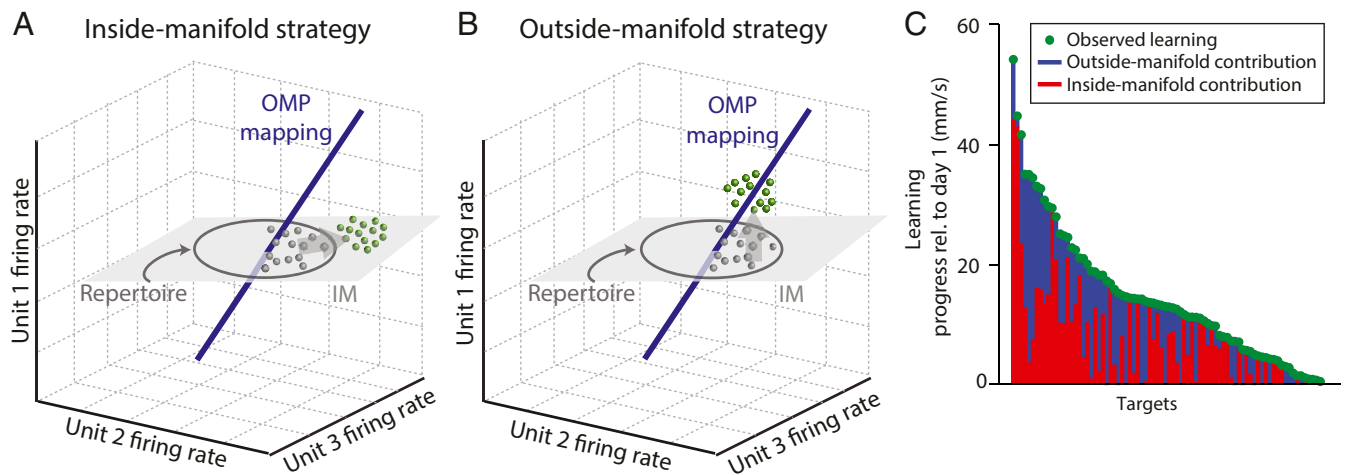


Fig. 4. Monkeys can produce new neural activity patterns outside of the intrinsic manifold. There are two types of new neural activity patterns: (A) those that are outside the repertoire, but remain within the manifold, and (B) those that are outside the manifold. Either type can yield performance improvements. (C) Animals learn using both inside-manifold and outside-manifold strategies for a given OMP mapping. Each bar shows one target from one multiday OMP experiment. The overall learning, defined as change in progress from day 1 to the last day, is represented by the green dot. Targets are ordered based on the amount of learning. The inside-manifold contributions to that learning are shown in red. The outside-manifold contributions are shown in blue. For visual clarity, data presented here show only targets with behavioral improvement and only the helpful contributions are shown. See *SI Appendix, Fig. S6* for a full presentation of these data.

to progress (*SI Appendix, Materials and Methods*). The new neural activity patterns on the last day of a multiday learning experiment include patterns with both substantial inside- and outside-manifold contributions to progress (Fig. 4C and *SI Appendix, Fig. S6 C–E*). This means that animals learned to move the cursor to some targets by generating new neural activity patterns that were outside of the intuitive neural repertoire but within the manifold, whereas other targets were learned by generating new neural activity patterns that were both outside of the intuitive neural repertoire and outside of the intrinsic manifold. Thus, learning can occur by changing the correlation structure (*SI Appendix, Fig. S6C*, blue), and also by changing firing rates in a manner that preserves the correlation structure (*SI Appendix, Fig. S6C*, red). We conclude that the brain can overcome the neural constraints imposed by the intrinsic manifold that we previously observed during single-day learning (8), but doing so takes several days.

Discussion

We found support for our hypothesis that populations of neurons can produce new patterns of activity to enable new behavioral capacities. Our BCI approach allowed us to establish a causal link from changes in neural activity patterns to changes in behavior. We encouraged new neural activity patterns to form by presenting the monkey with novel BCI mappings (OMPs). We detected the emergence of new patterns of neural activity after several days of practice with each OMP mapping. These patterns led directly to improved behavior. Some new patterns conformed to the preexisting correlation structure among the population of neurons, whereas other new patterns represented changes in the correlation structure. We expect that skill learning in general may proceed in part by the formation of new neural activity patterns appropriate for the learned behavior (19).

We can interpret changes in the brain at other levels of organization in the context of changes in neural populations. For example, learning has been associated with synaptic plasticity and the expansion of cortical maps. The time course of learning that we observed is consistent with findings that cortical synaptogenesis and motor map reorganization occur during late phases of skill learning (20). We posit that OMP learning involves synaptic plasticity. In fact, a function of synaptic plasticity

could be to permit a network of neurons to generate new patterns of activity. This plasticity may occur among the neurons in M1 from which we record, or among neurons from which we do not record, but that drive the recorded population.

Learning has been associated with changes in tuning properties of individual neurons in previous BCI (6, 7) and motor skill learning (21) studies. We posit that the formation of new neural activity patterns during long-term BCI learning may provide a parsimonious explanation for the tuning curve changes reported in earlier studies. In particular, our results combined with earlier BCI studies (6–9, 15, 16, 18, 22–26) and motor learning studies (27, 28) suggest that fast and slow learning are driven by different neural mechanisms. Fast learning can be accomplished by reassociating preexisting patterns of neural activity with new behaviors (9). This would result in neural tuning changes that are coordinated across the population (6, 29). Slow learning, as reported here for OMP mappings, can involve the formation of new patterns of neural activity. This would correspond to neural tuning changes that are specific to individual neurons (29). Other slowly learned phenomena that may also entail the creation of new patterns of neural population activity include motor skill, or de novo (30), learning (31, 32), cognitive learning (33, 34), and stroke recovery (35, 36). In the future, it might be possible to facilitate learning in those contexts by directly guiding the formation of new neural activity patterns appropriate for the desired behavioral capacities, as we have done here.

Materials and Methods

We used a BCI paradigm in which the monkey moved a computer cursor from the center of a screen to one of eight peripheral targets by volitionally modulating the activity of a population of ~90 neural units recorded with a multielectrode array chronically implanted in the arm region of primary motor cortex. Each experiment began by presenting the monkey with an “intuitive” mapping that relates neural activity to cursor velocities in a way that provides proficient control without requiring the animal to learn. Then, we induced learning by presenting the monkey with a novel mapping from neural activity to cursor velocity. The novel mappings encouraged the formation of new neural activity patterns. Subsequent analyses, demonstrate that the novel mappings are learned by the formation of new patterns of neural activity, and that those patterns directly drive the new behavior.

ACKNOWLEDGMENTS. We thank Patrick Sadtler for the original single-day learning results that inspired this work and for data included here for comparison (monkey L's single-day WMP and single-day no-incremental-training OMP data). This work was funded by NIH R01 HD071686 (to A.P.B., B.M.Y., and S.M.C.), National Science Foundation (NSF) BCS1533672 (S.M.C., B.M.Y., and A.P.B.), the Burroughs Wellcome Fund (A.P.B.), NSF

CAREER Award IOS1553252 (to S.M.C.), NIH CRCNS R01 NS105318 (to B.M.Y. and A.P.B.), NIH R01 HD090125 (to A.P.B.), Craig H. Neilsen Foundation 280028 (to B.M.Y., S.M.C., and A.P.B.), Pennsylvania Department of Health Research Formula Grant SAP 4100077048 under the Commonwealth Universal Research Enhancement program (to S.M.C. and B.M.Y.), Simons Foundation 543065 (to B.M.Y.), and NIH T32 NS07391 (to E.R.O. and A.D.D.).

1. T. V. Bliss, T. Lomo, Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* **232**, 331–356 (1973).
2. G. Q. Bi, M. M. Poo, Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* **18**, 10464–10472 (1998).
3. R. J. Nudo, G. W. Milliken, W. M. Jenkins, M. M. Merzenich, Use-dependent alterations of movement representations in primary motor cortex of adult squirrel monkeys. *J. Neurosci.* **16**, 785–807 (1996).
4. C. Gaser, G. Schlaug, Brain structures differ between musicians and non-musicians. *J. Neurosci.* **23**, 9240–9245 (2003).
5. M. D. Golub, S. M. Chase, A. P. Batista, B. M. Yu, Brain-computer interfaces for dissecting cognitive processes underlying sensorimotor control. *Curr. Opin. Neurobiol.* **37**, 53–58 (2016).
6. B. Jarosiewicz *et al.*, Functional network reorganization during learning in a brain-computer interface paradigm. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 19486–19491 (2008).
7. K. Ganguly, J. M. Carmena, Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biol.* **7**, e1000153 (2009).
8. P. T. Sadtler *et al.*, Neural constraints on learning. *Nature* **512**, 423–426 (2014).
9. M. D. Golub *et al.*, Learning by neural reassociation. *Nat. Neurosci.* **21**, 607–616 (2018).
10. B. A. Linkenhoker, E. I. Knudsen, Incremental training increases the plasticity of the auditory space map in adult barn owls. *Nature* **419**, 293–296 (2002).
11. S.-W. Park, T. M. A. Dijkstra, D. Sternad, Learning to never forget: Time scales and specificity of long-term memory of a motor skill. *Front. Comput. Neurosci.* **7**, 111 (2013).
12. J. E. Downey, N. Schwed, S. M. Chase, A. B. Schwartz, J. L. Collinger, Intracortical recording stability in human brain-computer interface users. *J. Neural Eng.* **15**, 046016 (2018).
13. A. S. Dickey, A. Suminski, Y. Amit, N. G. Hatsopoulos, Single-unit stability using chronically implanted multielectrode arrays. *J. Neurophysiol.* **102**, 1331–1339 (2009).
14. R. D. Flint, M. R. Scheid, Z. A. Wright, S. A. Solla, M. W. Slutzky, Long-term stability of motor cortical activity: Implications for brain machine interfaces and optimal feedback control. *J. Neurosci.* **36**, 3623–3632 (2016).
15. J. A. Hennig *et al.*, Constraints on neural redundancy. *eLife* **7**, e36774 (2018).
16. E. J. Hwang, P. M. Bailey, R. A. Andersen, Volitional control of neural activity relies on the natural motor repertoire. *Curr. Biol.* **23**, 353–361 (2013).
17. A. Luczak, P. Barthó, K. D. Harris, Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron* **62**, 413–425 (2009).
18. S. Sakellaridi *et al.*, Intrinsic variable learning for brain-machine interface control by human anterior intraparietal cortex. *Neuron* **102**, 694–705.e3 (2019).
19. H. Makino, E. J. Hwang, N. G. Hedrick, T. Komiyama, Circuit mechanisms of sensorimotor learning. *Neuron* **92**, 705–721 (2016).
20. J. A. Kleim *et al.*, Cortical synaptogenesis and motor map reorganization occur during late, but not early, phase of motor skill learning. *J. Neurosci.* **24**, 628–633 (2004).
21. C. S. Li, C. Padoa-Schioppa, E. Bizzi, Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. *Neuron* **30**, 593–607 (2001).
22. A. J. Law, G. Rivlis, M. H. Schieber, Rapid acquisition of novel interface control by small ensembles of arbitrarily selected primary motor cortex neurons. *J. Neurophysiol.* **112**, 1528–1548 (2014).
23. V. R. Athalye, K. Ganguly, R. M. Costa, J. M. Carmena, Emergence of coordinated neural dynamics underlies neuroprosthetic learning and skillful control. *Neuron* **93**, 955–970.e5 (2017).
24. M. G. Perich, J. A. Gallego, L. E. Miller, A neural population mechanism for rapid learning. *Neuron* **100**, 964–976.e7 (2018).
25. V. R. Athalye, F. J. Santos, J. M. Carmena, R. M. Costa, Evidence for a neural law of effect. *Science* **359**, 1024–1029 (2018).
26. S. Vyas *et al.*, Neural population dynamics underlying motor learning transfer. *Neuron* **97**, 1177–1186.e3 (2018).
27. D. J. Berger, R. Gentner, T. Edmunds, D. K. Pai, A. d'Avella, Differences in adaptation rates after virtual surgeries provide direct evidence for modularity. *J. Neurosci.* **33**, 12384–12394 (2013).
28. R. Shadmehr, M. A. Smith, J. W. Krakauer, Error correction, sensory prediction, and adaptation in motor control. *Annu. Rev. Neurosci.* **33**, 89–108 (2010).
29. X. Zhou, R. N. Tien, S. Ravikumar, S. M. Chase, Distinct types of neural reorganization during long-term learning. *J. Neurophysiol.* **121**, 1329–1341 (2019).
30. J. W. Krakauer, A. M. Hadjiosif, J. Xu, A. L. Wong, A. M. Haith, Motor learning. *Compr. Physiol.* **9**, 613–663 (2019).
31. D. Sternad, It's not (only) the mean that matters: Variability, noise and exploration in skill learning. *Curr. Opin. Behav. Sci.* **20**, 183–195 (2018).
32. L. Shmuelof, J. W. Krakauer, P. Mazzoni, How is a motor skill learned? Change and invariance at the levels of task success and trajectory control. *J. Neurophysiol.* **108**, 578–594 (2012).
33. Y. Gu *et al.*, Perceptual learning reduces interneuronal correlations in macaque visual cortex. *Neuron* **71**, 750–761 (2011).
34. J. M. Jeanne, T. O. Sharpee, T. Q. Gentner, Associative learning enhances population coding by inverting interneuronal correlation patterns. *Neuron* **78**, 352–363 (2013).
35. S. R. Soekadar, N. Birbaumer, M. W. Slutzky, L. G. Cohen, Brain-machine interfaces in neurorehabilitation of stroke. *Neurobiol. Dis.* **83**, 172–179 (2015).
36. R. J. Nudo, Mechanisms for recovery of motor function following cortical damage. *Curr. Opin. Neurobiol.* **16**, 638–644 (2006).