

SENSORIMOTOR LEARNING

Set in one's thoughts

A key component of learning involves updating existing motor plans in response to altered sensory feedback. By using a brain–computer interface, Golub et al. show how such learning changes the activity of neural populations in primary motor cortex—and how it does not.

Aniruddh R. Galgali and Valerio Mante

Before taking a leap through a tree, a monkey must estimate the strength of the next branch and the distance to it, plan the jump and then execute it. The nature of this sensory–motor transformation, from visual inputs to movement commands, is subject to learning, as it must be continually adjusted to account for changes in the body (growth) or the environment (a kind of tree never encountered before). Sometimes, if the animal is to survive, adjustments must occur fast. For example, if a monkey is injured and temporarily cannot use one of its limbs, it must rapidly learn to produce new movements that do not rely on that limb.

Probing the neural correlates of such sensorimotor learning is challenging. At least in primates, typically only a small subset of the neurons whose activity might be changing through learning can be experimentally monitored, and a causal link between observed learning-related changes in neural activity and changes in motor behavior is often difficult to prove.

These challenges can be circumvented with brain–computer interfaces¹, whereby the movement of an artificial effector (for example, a cursor on a screen) is controlled directly through the activity of a set of recorded neurons. As the mapping between neural activity and effector movements is chosen by the experimenter, brain–computer interfaces provide complete experimental control both over which neurons are causally responsible for the observed behavior and over the nature and temporal dependence of the sensory–motor mapping.

In this issue of *Nature Neuroscience*, Golub et al.² use a brain–computer interface to study how neural activity in primary motor cortex (M1) changes as monkeys adapt to a sudden perturbation in the mapping that transforms M1 activity directly into the movement of a cursor on a screen. Using a clever experimental design and sophisticated computational methods, the authors predict how neural activity would have to change through learning to allow the monkeys to perform the task optimally. They find that, while the monkeys can adjust to a

perturbation quite effectively, they fall short of producing the changes in neural activity that would be required for optimal behavior. The reason for this shortcoming appears to be a striking inability to rapidly change the overall repertoire of patterns of population activity produced in M1. Instead, learning occurs by re-associating the activity patterns already present in the fixed repertoire with different cursor movements after the perturbation.

The goal of the monkeys was to move a cursor from a central point on the screen to a peripheral visual target, whose location varied from trial to trial (Fig. 1). At any given time during the trial, the cursor velocity was directly controlled by the pattern of activity simultaneously recorded in M1 by means of a multielectrode array. The nature of this readout is best understood in a state-space representation of the neural population activity (Fig. 1a) wherein each data point corresponds to the pattern of population activity measured at a given time for a given target location. The cursor velocity along the x and y directions is then obtained by ‘projecting’ the population activity at each time onto two directions, V_x and V_y , in this state space (Fig. 1b). When these readout directions are kept fixed across days, monkeys can learn over the course of weeks or months to generate patterns of activity (Fig. 1a) that result in the appropriate cursor movements for a given target (Fig. 1c).

The critical experimental manipulation was to change the readout directions midway through an experimental session, from an ‘intuitive’ set of directions that the monkeys had already mastered to an entirely novel, never-before-seen set of ‘perturbed’ directions (Fig. 1b). Right after the perturbation, the monkeys still generate patterns that would be appropriate for the intuitive readout but that do not produce the desired cursor movements under the perturbed readout (Fig. 1c). To reduce the resulting mismatch between the desired and observed cursor movements, the monkeys thus had to adapt to the perturbation by changing the neural activity generated in M1 for any given target.

The key finding of Golub et al.² is that, at least in the course of the 1–2 h following

a perturbation, the monkeys appear unable to adapt to the perturbation by applying the strategy that would result in the most accurate and fastest cursor movements. Given the nature of the perturbations chosen by Golub et al.², this optimal strategy would require ‘realigning’ the overall repertoire of population patterns to achieve the same relative arrangement of population patterns and readout directions as before the perturbation (Fig. 1d). Rather than realigning the responses, the monkeys instead seem to follow the suboptimal strategy of ‘re-association’: the patterns that already exist in the original neural repertoire are reused, but they are generated for different target locations than before the perturbation (Fig. 1d). Providing evidence for re-association and against realignment is very much nontrivial, and indeed much of the effort by Golub et al.² goes into developing the necessary computational and analytical tools.

Notably, realignment and related strategies such as rescaling (see ref. 2) require the generation of novel population patterns that were not part of the original neural repertoire produced in the intuitive mapping (Fig. 1d). Unlike in a previous study³ by the same authors, which showed an inability to adapt to very strong perturbations, here the perturbations are less severe in that they lie within the ‘intrinsic manifold’ spanned by the shared covariance in the neural population^{2,3}. As a result, here the required novel patterns could have been produced simply by increasing the strength of certain patterns in the original repertoire (i.e., by increasing their distance from the origin of the state space in Fig. 1a while maintaining the relative projections onto the state-space axes). By showing that monkeys cannot even change the strength of already present patterns, Golub et al.² demonstrate that dynamics in areas like motor cortex are much less malleable than one might have thought.

One important assumption underlying the authors’ conclusions is that the monkeys’ inability to implement the optimal strategy is not simply due to a lack of motivation. As the authors themselves acknowledge, one cannot completely exclude the possibility

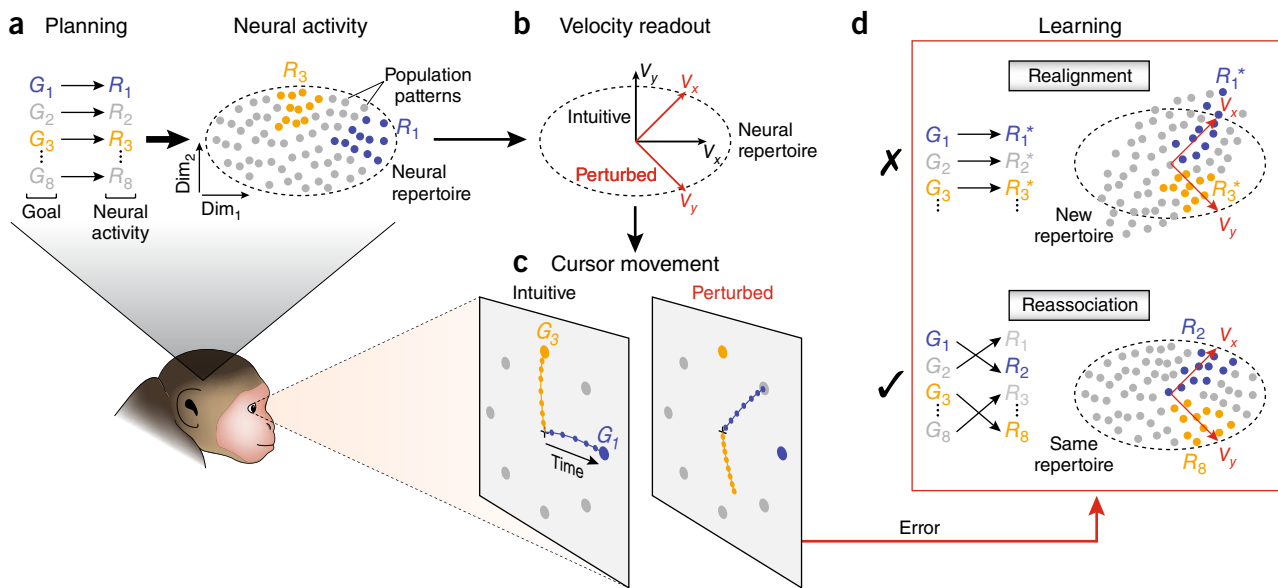


Fig. 1 | Neural correlates of learning in a brain-machine interface experiment. A monkey is trained to move a cursor toward a goal target on a video monitor. The velocity of the cursor is controlled directly by the momentary neural activity in primary motor cortex (M1). Midway through an experiment, the mapping between neural activity and cursor velocity is perturbed, forcing the monkey to relearn how to control the cursor. **a**, Before a perturbation, each goal generates M1 activity resulting in appropriate cursor movements (planning). The observed population patterns in M1 are represented as points in a low-dimensional state space and span the neural repertoire for this task. Dim, dimension. **b**, Each pattern is projected onto readout directions V_x and V_y to obtain the corresponding horizontal and vertical cursor velocities. A perturbation is achieved by switching from an intuitive (i.e., previously mastered; black) readout to two novel directions (red) in a ten-dimensional state-space (two dimensions illustrated). **c**, Cursor movements are accurate before (left) but not immediately after the perturbation (right). **d**, Over short time scales, learning is achieved by changing the mapping between goals and neural responses while the neural repertoire remains unchanged (re-association). Realignment of the neural repertoire, which results in optimal cursor movements but would require the generation of novel population patterns, is not observed.

that the monkeys could have changed the neural repertoire if they had ‘tried harder’. One of the observations speaking against this interpretation is that, among the many possible re-association solutions for a given perturbation, monkeys chose those that led to the most accurate cursor movements. In this sense, at least within the constraints of re-association, the monkeys do perform optimally. It thus seems plausible to conclude that when the mapping from a desired goal to an intended movement needs to be changed quickly, over the course of minutes to hours, the monkeys must do their best with the repertoire of neural activity that they already have, resulting in re-association (Fig. 1d). The addition of new patterns to the repertoire, and resulting realignment, may still be possible, but only over the course of days or weeks.

The reasons why learning occurs in this way might be found by taking a dynamical systems view of motor preparation and execution⁴. In this framework, preparatory activity in premotor and motor areas is thought to encode the parameters of an upcoming movement in the form of the initial condition of a dynamical system. The initial conditions may reflect the state of an input into motor cortex⁵. The complex sequences of activity patterns observed during movement can instead be interpreted as reflecting the

recurrent dynamics in motor cortex⁵. When the initial conditions are varied, the resulting sequences of activity patterns also change, producing different movements. Re-association between neural activity patterns and cursor movements could be brought about by merely changing the mapping between the movements and the already used initial conditions, while realignment would require a change in the rules governing the recurrent dynamics. From a machine-learning standpoint it then seems plausible that re-association would be easier to learn than realignment. Recurrent networks are notoriously hard to train⁶, as the necessary error signals are difficult to back-propagate through the recurrent connectivity. This ‘credit assignment’ problem, however, would not apply to an input specifying the initial conditions.

Further experiments and modeling will be required to determine whether and how activity in premotor and motor areas maps onto these putative inputs and recurrent dynamics, as well as how changes in these processes ultimately relate to fast and slow learning in more natural settings⁸. Even in the setting of brain-computer interfaces, it is not immediately clear how the findings of Golub et al.², which focus entirely on population-level analyses, relate to a substantial body of research that has identified a variety of changes in the activity of single units during learning^{9,10}.

However, the conceptual clarity and theoretical rigor in the arguments of Golub et al.² provide the basis for relating changes at the single-unit and population levels and set a solid foundation for further advances in deciphering the neural mechanisms of learning. □

Aniruddh R. Galgali* and Valerio Mante*
Institute of Neuroinformatics, University of Zurich
and ETH Zurich, Zurich, Switzerland.
*e-mail: galgalia@ini.uzh.ch; valerio@ini.uzh.ch

Published online: 12 March 2018
<https://doi.org/10.1038/s41593-018-0105-5>

References

1. Fetz, E. E. *J. Physiol. (Lond.)* **579**, 571–579 (2007).
2. Golub, M. D. et al. *Nat. Neurosci.* <https://doi.org/10.1038/s41593-018-0095-3> (2018).
3. Sadtler, P. T. et al. *Nature* **512**, 423–426 (2014).
4. Shenoy, K. V., Sahani, M. & Churchland, M. M. *Annu. Rev. Neurosci.* **36**, 337–359 (2013).
5. Sussillo, D., Churchland, M. M., Kaufman, M. T. & Shenoy, K. V. *Nat. Neurosci.* **18**, 1025–1033 (2015).
6. Bengio, Y., Simard, P. & Frasconi, P. *IEEE Trans. Neural Netw.* **5**, 157–166 (1994).
7. Athalye, V. R., Ganguly, K., Costa, R. M. & Carmena, J. M. *Neuron* **93**, 955–970.e5 (2017).
8. Paz, R., Boraud, T., Natan, C., Bergman, H. & Vaadia, E. *Nat. Neurosci.* **6**, 882–890 (2003).
9. Jarosiewicz, B. et al. *Proc. Natl. Acad. Sci. USA* **105**, 19486–19491 (2008).
10. Ganguly, K. & Carmena, J. M. *PLoS Biol.* **7**, e1000153 (2009).

Competing interests

The authors declare no competing financial interests.