



Learning outside the box

Raviv Pryluk^a and Rony Paz^{a,1}

Learning a new skill, whether it is riding a bicycle or playing chess, usually requires at least several days of practice if not much more. Due to experimental limitations, however, most neuroscience studies that attempt to unveil the neural representations of skill learning explore tasks that can be learned within much shorter time scales, a few hours at most. Therefore, the neural substrates of long-term learning remain poorly understood. An intriguing important question is why some tasks can be learned within hours, whereas other tasks require longer-term practice. Oby et al. (1) tackle this challenge by using a brain–computer interface (BCI) approach combined with chronic neural recordings. Although, originally, BCI was used to demonstrate the feasibility of designing neural prostheses (2, 3), a clever exploitation of this technique can help to discover properties of the neural code and how it changes during learning. In this study, they show that new patterns of neural activity emerge with long-term learning—patterns that were not previously observed in the neural population they measured. Some of these new patterns enable the longer-term learning.

A common paradigm to study motor control is the arm reaching paradigm in which primates move a manipulandum to control the location of a cursor on the screen, while neurons are recorded in motor regions of the brain. Although this paradigm has led to many successful discoveries about how neurons represent directional movements and new motor skills (4, 5), it remains unclear why some skills can be acquired within minutes or hours, whereas other skills require a much lengthier learning process that lasts days or weeks. A BCI paradigm, as used here, is one that directly links the changes in the monkey brain, namely the patterns of neural activity, with the desired outcome—the moving of a cursor to a desired location, in this particular case (Fig. 1A). In general, it allows defining a specific mapping function between neural activity and behavior so that it addresses the question at hand. Here, the main goal is to ask which mapping functions are easier or harder to learn, and, as a result, one can determine which neural patterns can be attained faster

and drive learning, and which require more practice and emerge later.

The BCI is fed with real-time activity of ~90 neurons recorded in the arm region of the primary motor cortex, a region that is known to show directional- and velocity-related responses during movement generation. This means that, in each time point, the combined activity is actually a vector in a 90-dimensional (90D) space. By using a dimensionality reduction technique, Oby et al. (1) first noticed that a much lower number of dimensions (10D) is sufficient to explain the observed neural patterns. This relatively low number of dimensions was chosen such that it explains most of the variance in the 90D neural patterns that occur during movement (more than 95%). The fact that 10 out of 90 dimensions explain so much already hints that the actual neural code potentially lies in a low-dimensional space, termed a manifold, or an intrinsic manifold (IM). Importantly, it also makes it possible to map the neural activity into the 2D movement of the cursor, so that, for example, if the neural activity resides in the left part of this manifold, the cursor will move left, and, if it resides in the right side, it will move right. In the real experiment, the mapping was a bit more complex between the 10D vector and 2D cursor velocity, yet the concept is similar (Fig. 1 A and B).

Using this sophisticated approach, this group has previously shown that short-term learning can be achieved when the mapping is constrained to a “within-manifold perturbation” (WMP; Fig. 1C) (6). When the animals are required to learn by using patterns of neural activity which lie on this low-dimensional manifold of neural patterns that were characterized prelearning, they can do so reasonably over a few continuous hours within 1 d. Moreover, it was shown that changes in population activity after such short-term learning followed a strategy they term “reassociation” (Fig. 1C). This strategy uses the existing repertoire of activity patterns, yet associates them differently with movements after learning (7). As a simplified example, if patterns A, B, C, and D occurred during movement before learning (dots on the

^aDepartment of Neurobiology, Weizmann Institute of Science, Rehovot 76100, Israel

Author contributions: R. Pryluk and R. Paz wrote the paper.

The authors declare no conflict of interest.

Published under the [PNAS license](#).

See companion article 10.1073/pnas.1820296116.

¹To whom correspondence may be addressed. Email: rony.paz@weizmann.ac.il.

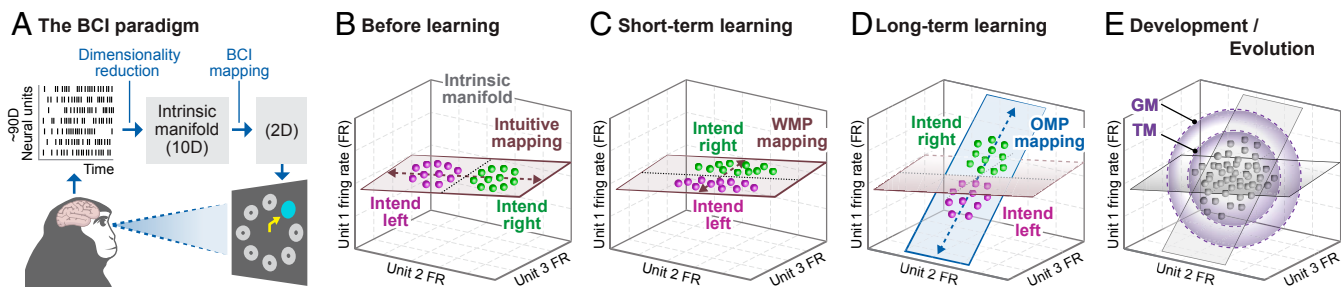


Fig. 1. Long-term learning requires emergence of new neural patterns: a BCI approach. (A) Schematic of the BCI paradigm: Monkeys controlled a cursor (yellow) to acquire 1 (cyan) of 8 possible (gray) targets by modulating their own neural activity. The activity of ~90 neurons in the motor cortex is recorded, and this 90D activity space is reduced into a 10D IM by a dimensionality reduction technique, and then mapped to 2D cursor velocity. **(B)** A simplified conceptual schematic of observed—recorded neural activity patterns (dots). The neural activity tends to lie in a low-dimensional subspace, termed the IM (gray plane). Monkeys move the BCI cursor by volitionally modulating their neural activity. Initially, cursor velocities were determined by the previously observed patterns (termed “intuitive” mapping, black arrow). Here, green patterns would result in cursor movement to the right, and purple patterns would result in left movement. **(C)** To induce short-term learning, WMP mapping is imposed. Under the WMP, neural activity patterns map to different cursor velocities than under the intuitive mapping. This encourages the monkeys to learn, and it was shown that they can do so within a few hours, mainly by reassociation of existing patterns to different outcome results, i.e., directional cursor movement. **(D)** The same as in C, but for OMP. This entailed much longer learning, lasting several days, and resulted in emergence of new activity patterns that were not observed before. **(E)** TM is a manifold that bounds all of the activity patterns observed during a specific task. However, it is highly likely that there is a GM, which bounds this specific TM as well as other TMs for different tasks. Such a GM would be constrained by the efficiency and robustness of the neural code, determined by network interactions and architecture. How to unveil this GM is a future goal that would yield important insights on how the brain enables as well as constraints adaptive as well as maladaptive behaviors.

manifold in Fig. 1B), where A and B moved the cursor leftward and C and D moved it rightward, the monkey could learn to produce (reassociate) A, C to move leftward and B, D to move rightward, when such mapping was imposed (e.g., dots in Fig. 1C).

In contrast to WMP, “outside-manifold perturbations” (OMP; Fig. 1D) are not well learned within a single day (6). OMP means that the monkeys’ neurons are required to generate patterns of activity that are novel and were not observed before (or observed only rarely, e.g., patterns E, F). Oby et al. (1) leverage this to induce long-term learning and to study its neural mechanism. They waited patiently to allow the monkeys to train day after day, a real “tour de force” in such hard monkey experiments, where most studies stop after a few days if no learning occurs. When they noticed that learning was still too hard, they cleverly facilitated it by using an incremental training strategy (8). This indeed allowed learning to proceed with gradual improvements from day to day, albeit with dips and rebounds in the performance that likely reflect the difficulty of acquiring the new skill. Importantly, the chronic recordings of the same neurons across days enabled detection of new activity patterns that did not exist before learning. The evidence is clear. First, the number of emerging new patterns is correlated with the progress in learning performance, and second, there are 2 types of new patterns: (i) patterns that reside on the outside-manifold they imposed; and (ii) patterns that reside on the IM but that can contribute to the new mapping.

Together, the results nicely show that the brain can learn to generate new activity patterns to achieve the desired goal, but it requires more time if these patterns are very different from the patterns that already existed and were used before learning started (9). This provides a simple intuition and potential explanation why some tasks are harder to learn and take more practice time—because the brain is required to learn to produce new patterns of neural activity. Moreover, it is not only new patterns but new patterns that are outside of the IM, meaning that formerly inactive neurons become activated. In future studies, it might be important to expand and understand more precisely the contribution of completely new patterns and the contribution by

reassignment of old patterns (similar to the reassociation they previously showed).

As an analogy to the finding, someone who is an experienced surfer will likely easily learn to snowboard, yet it would require much more time to learn to play soccer well. This is despite the fact that all 3 activities require mainly mid-to-lower body motor skills, and because they require different muscle and joint combinations which are somewhat more similar for snowboard and wave surfing than for soccer. If neural activity in primary motor cortex is related to generation of muscle–joint combinations, and not only to direction of movement and its velocity, then perhaps the manifold identified before learning (but after substantial training on the 2D task) relates to similar movement on 2 dimensions, and therefore new mapping that is based on WMP mainly means mapping of velocity and direction, namely, activating the same muscle–joint pairs but in different combinations of intensity. This might resemble more motor “adaptation,” either a visuomotor or a force-field one, and might be identical to swimming in water with different viscosities/densities, or running with or against the wind. In contrast, a mapping that requires OMP patterns might be more similar to the performance of new joint–muscle activations that were rarely performed before, as if learning a totally new motor skill.

Here, the space of neural activity patterns was defined during same-task performance; this is, of course, a clever approach that likely helped obtain the current results. It also raises the question of what would happen if they characterized the baseline activity and used it to define the IM. Would it be a much more complex one, i.e., requiring many more dimensions to reach 95% explained variance? If so, would it mean that OMP would be even harder to learn, as new patterns would have to come from an even less likely repertoire? Alternatively, it could well be that baseline ongoing activity characterizes network states and that stimulus- or movement-evoked patterns are a subset of these patterns (10). In this case, identifying the IM would allow designing a richer set of tasks that lie on the WMP and can be learned relatively fast.

