Brain–computer interfaces

https://doi.org/10.1038/s41593-024-01793-2

Neural populations are dynamic but constrained

Amy L. Orsborn

Our brains evolved to help us rapidly learn new things. But anyone who has put in hours of practice to perfect their tennis serve, only to reach a plateau, can attest that our brains aren't infinitely flexible. New work shows that patterns of neural activity over time – the temporal dynamics of neural populations – cannot change rapidly, suggesting that neural activity dynamics may both reflect and constrain how the brain performs computations.

Neurons are the fundamental unit within the brain, but neurons work together to control our behavior. As technology made it easier to measure the activity of many neurons simultaneously, neuroscientists noticed structure within their activity (Fig. 1a). For example, two neighboring neurons may have closely related activity, potentially because they receive similar inputs, perform similar computations, or share information with each other through a synaptic connection. There's also structure in how neural activity unfolds in time (Fig. 1b). These temporal patterns reflect the fact that neural circuits are physical systems with inherent dynamics, which are determined in part by the inputs to and connections between neurons. Past work shows that population dynamics correlate with behavioral variables, such as planning and executing movements[1]. Analysis of dynamics within artificial neural network (ANN) models also helps to reveal how nodes in the network collaborate to perform computations[2].

While evidence hints that the dynamics of neural populations reflect the underlying computations of the brain, testing this hypothesis brings up a common question in neuroscience: are the patterns we observe circumstantial correlations, or do they capture fundamental properties of how the brain always operates? Oby, Degenhart, Grigsby and colleagues used brain–computer interfaces (BCIs) to provide the most direct evidence to date that the activity of neurons in the motor cortex may be constrained to follow certain temporal dynamics[3]. BCIs create experimenter-defined sensory–motor systems in which neural activity is mapped into movement outputs (Fig. 1c). BCIs are best known as a therapy to restore movements to people with paralysis[4], but this study adds to a body of work that uses BCIs to precisely probe how neural activity gives rise to behaviors[5,6]. The authors used the ability to manipulate how neural activity relates to movement to directly test whether their subjects, rhesus macaques, could readily alter the temporal dynamics of activity in motor cortex. If not, they reasoned, this suggests that temporal patterns reflect meaningful properties of brain circuits.

The study relied on an innovative method that, in effect, made the dynamics of the brain directly visible to the monkeys via the BCI (Fig. 1c). They trained the monkeys to imagine moving a cursor left and right on a screen without making actual movements. Monkeys, like humans, aim to do this task in the most efficient way possible – moving in a straight line. The authors measured the activity of about 100 neurons while monkeys imagined movements, and were able to summarize this neuronal activity to a combination of 10 patterns, reflecting neurons working together. Within this 10-dimensional neural space, they found a 2D projection that corresponded to the cursor trajectories going straight back and forth between the targets. This mapping is what typical BCIs would use, as it is thought to capture the monkeys' intended movements. But Oby, Degenhart, Grigsby et al. discovered an additional, different 2D projection of the same neural activity patterns that produced trajectories that hit the left and right targets, but with curved instead of linear paths. This curvature stemmed from the fact that the temporal patterns of neural activity – the exact sequence of how neurons were activated in time – were different when the monkeys imagined movements to the left or the right. The researchers showed the monkeys this alternate view of the neural activity patterns, with the temporal differences visible. When they used this BCI mapping, the monkeys would imagine moving straight as before, but would see the cursor trace curved paths (Fig. 1d).

By creating a BCI mapping in which temporal dynamics influenced the cursor's position, the researchers could now test whether the monkeys could alter the temporal dynamics of motor cortex so that the cursor paths would become straight. They did this with a series of tasks each day that gradually increased the need for the monkeys to change the dynamics to solve the task. They first investigated whether the monkeys might alter the temporal dynamics on their own, as we know that monkeys prefer to take the most efficient, straight path. After several hundred repetitions, the monkeys kept making curved movements. After those initial trials, they modified the task to more explicitly incentivize the monkeys to change the temporal dynamics, including a task that required monkeys to modify their dynamics to get rewards. The monkeys struggled with these tasks despite these extra incentives. While it is fundamentally impossible to prove that monkeys can't alter temporal dynamics, the experiments show that the monkeys didn't do so within the timescales of the experiment and despite incentives.

The unique strengths of BCI allowed the researchers to demonstrate that the dynamics governing neural activity patterns are not readily changeable, which gives us new clues about how the brain performs computations. This finding provides compelling evidence that the spatiotemporal structure we observe in neural population recordings is likely to reflect meaningful properties of neural circuits. This bolsters our confidence that analyzing the dynamics of neural population activity will be beneficial for inferring computations performed by groups of neurons.

What exactly do neural population dynamics reflect within neural circuits? In ANN models, the temporal dynamics of nodes are determined in part by the connections between the nodes. The authors therefore argue that their results support the interpretation of neural activity as reflecting connections between neurons. Their findings

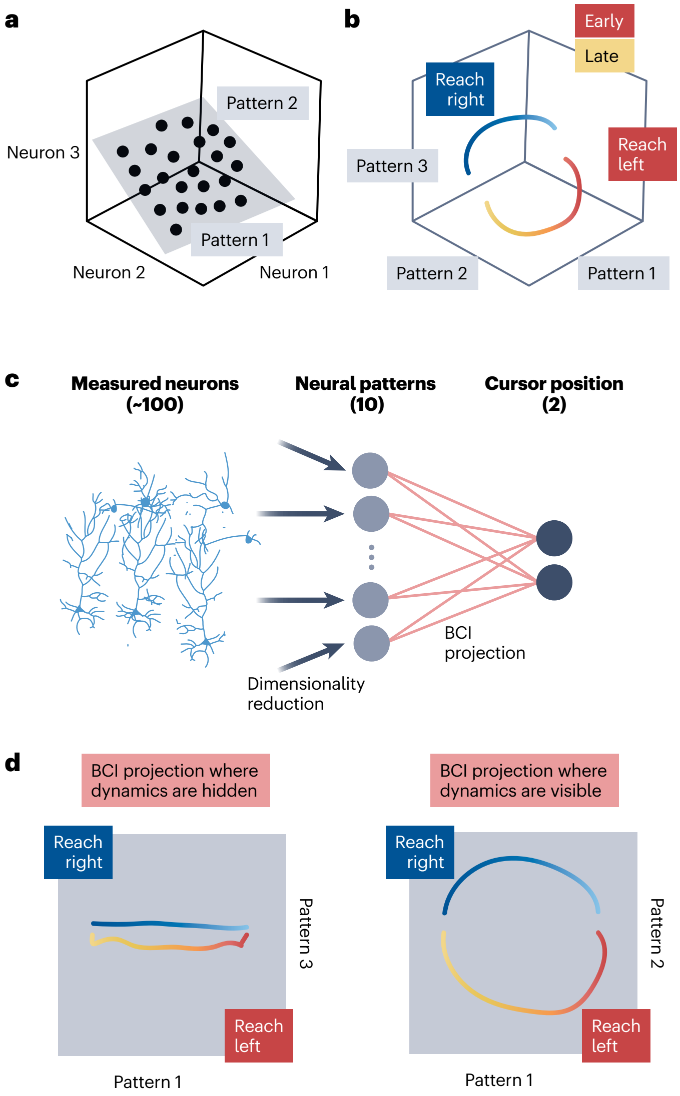

Fig. 1 | BCIs that make neural dynamics visible. Neural activity has rich spatiotemporal structure. **a**, Examining the activity of neurons relative to each other (black dots) shows that neuronal activity is correlated. This creates patterns of activity that are lower-dimensional than the total number of neurons (gray plane). **b**, Within the space of possible patterns among neurons, neural activity evolves in structured ways over time (trajectories). In the paper by Oby et al.[3], the temporal sequences of neural activity patterns were different when monkeys aimed to make movements to the right (blue) or to the left (red). **c**, A BCI mapped neural activity into the position of a cursor in two stages: 1) dimensionality reduction (middle), which transformed the activity of individual neurons (about 100) into the space of coordinated patterns (10); and 2) BCI projection (right), which mapped neural patterns (10) into cursor positions (2). **d**, Akin to looking at the plot in **b** from different viewing angles, the researchers altered the BCI projection to show the monkeys different views of their neural activity. One projection (left), which is typically used for BCI applications, hid the different temporal patterns between right and left movements from the monkeys. The other projection (right) showed the monkeys a view that maximized the difference between the two movements; here, the aim was to assess whether the monkeys could alter the temporal dynamics.

provide further confidence that temporal dynamics give us a lens into network connections, but open questions remain. Directly linking recorded neural activity patterns to physiological connections in the brain is challenging and error-prone, in part because we usually measure only a small portion of the network[7]. For instance, temporal dynamics in ANNs are also influenced by inputs to the network, which were not measured in these experiments. An interesting avenue for future inquiry will be to investigate how the glimpse of network structure we get from studying the temporal dynamics within one group of neurons relates to physiological connections.

This work also helps us to understand the contexts under which neural activity can and cannot change. Past BCI studies showed that animals can quickly learn to activate a small handful of neurons to solve tasks, suggesting vast flexibility[8]. This study, along with the researchers' past work[9], highlights that larger neural populations may not be able to change arbitrarily on short timescales. Interestingly, some forms of rapid motor learning are accomplished by changing the inputs to primary motor cortex[10], suggesting that the monkeys in this study could have used similar strategies for this task. The monkeys' inability to modify their temporal dynamics, then, suggests that there may be limits on how much modifying inputs to motor cortex can alter its temporal dynamics. Monitoring the activity of brain areas that project to motor cortex during these tasks will help to further tease apart what aspects of neural dynamics are and are not flexible.

This new observation raises fundamental questions about how and why the brain might constrain its dynamics. What physiological mechanisms maintain temporal dynamics within a network? There are likely to be interesting avenues to link these findings to research on homeostatic plasticity, sleep, and circadian mechanisms that regulate neural circuit functions. How does the brain generate our rich and adaptable behavior if our neural activity is constrained? This discovery hints that temporal dynamics within a network could provide a scaffold upon which learning happens. Why would our brains use this strategy? ANNs appear to re-use dynamics learned for one task to help them learn other tasks more quickly[2]. Perhaps our brains do something similar. Maybe there is a core neuroscientific principle underlying the common notion that creativity flourishes under constraints.

Amy L. Orsborn [1,2,3] ✉
[1]Department of Electrical & Computer Engineering, University of Washington, Seattle, WA, USA. [2]Department of Bioengineering, University of Washington, Seattle, WA, USA. [3]Washington National Primate Research Center, Seattle, WA, USA.
✉e-mail: aorsborn@uw.edu

References

1. Vyas, S., Golub, M. D., Sussillo, D. & Shenoy, K. V. *Annu. Rev. Neurosci.* **43**, 249–275 (2020).
2. Driscoll, L. N., Shenoy, K. & Sussillo, D. *Nat. Neurosci.* **27**, 1349–1363 (2024).
3. Oby, E. R. et al. *Nat Neurosci*. https://doi.org/10.1038/s41593-024-01845-7 (2025).
4. Pandarinath, C. & Bensmaia, S. J. *Physiol. Rev.* **102**, 551–604 (2022).
5. Orsborn, A. L. & Pesaran, B. *Curr. Opin. Neurobiol.* **46**, 76–83 (2017).
6. Golub, M. D., Chase, S. M., Batista, A. P. & Yu, B. M. *Curr. Opin. Neurobiol.* **37**, 53–58 (2016).
7. Das, A. & Fiete, I. R. *Nat. Neurosci.* **23**, 1286–1296 (2020).
8. Fetz, E. E. *J. Physiol. (Lond.)* **579**, 571–579 (2007).
9. Oby, E. R. et al. *Proc. Natl. Acad. Sci. USA* **116**, 15210–15215 (2019).
10. Perich, M. G., Gallego, J. A. & Miller, L. E. *Neuron* **100**, 964–976.e7 (2018).

Competing interests

A.L.O. is a scientific advisor for Meta Reality Labs.