Neural manifolds: more than the sum of their neurons

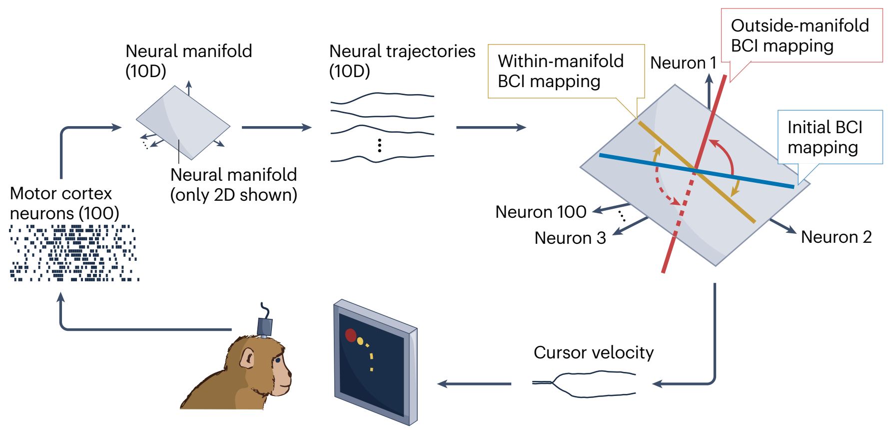

Together with lesion and stimulation studies, correlating the recorded activity of single neurons from behaving animals with relevant task variables allowed neuroscientists to draw the 'maps' of brain function found in textbooks. However, early in the 21st century, some researchers – inspired by the study of coordinated behaviour in complex systems – quietly initiated what became a small paradigm shift in systems neuroscience: moving the focus from single-neuron activity to the coordinated activity across neurons in a recorded population.

This transition seems to have begun in 2003, when Stopfer et al. realized that interpreting the complex single-neuron activity patterns evoked in the locust olfactory system by different odorants became easier when they focused on the neurons' coordinated activity. Using mathematical techniques to make sense of datasets comprising many variables – coordination of activity over time in a large population of more than a hundred neurons – they identified the main patterns of covariation across the recorded olfactory neurons. This allowed the authors to reduce the distinct activity patterns to three covariation patterns that captured the coordinated behaviour well. Plotting each of these main covariation patterns along an axis of a three-dimensional space revealed a beautiful arrangement: each odour evoked a similar-looking rotational trajectory that captured how the activity of the neural population evolved in response to that odorant. Moreover, all the trajectories evoked by the same odorant lived in the same plane, and increasing the intensity of an odour simply made these trajectories bigger. The – in this case – three-dimensional surface along which these trajectories evolved became known as a 'neural manifold'.

Probably inspired by this result's beauty, groups started to adopt this manifold approach to interpret the activity of the increasing neuron numbers they were recording – work that led to several fascinating findings that unfortunately cannot be covered here. But Sadtler et al. took this approach one step further by asking whether neural manifolds are more than a convenient way to interpret the population activity: what if they capture something fundamental about neural function?

Biological factors such as circuit connectivity constrain how neurons can fire with respect to each other, and despite its complexity this interaction can be loosely expressed as follows: the activity of closely connected neurons will covary. Thus, Sadtler et al. hypothesized that as neural manifolds capture the dominant patterns of covariation across neurons in a population, then they might reflect fundamental constraints on what activity patterns a neural population can produce. More specifically, the authors predicted that it should be much easier for an animal to produce population activity patterns with trajectories within an existing neural manifold than outside it, as the latter could potentially require new synaptic connections to allow those activity patterns.

To test this hypothesis, Sadtler et al. designed an elegant experiment in which they 'asked' macaque monkeys to generate new population activity patterns that were within or outside an existing manifold, which they identified at the beginning of each experimental session. To achieve this, they built a brain–computer interface (BCI) that gave two monkeys online 'mental control' of a computer cursor by mapping the trajectories described by the activity of about 100 of their motor cortex neurons within a ten-dimensional (10D) manifold onto the cursor velocity (see image). Then, after showing that monkeys could easily use this BCI to bring the cursor to a target to get a juice reward, Sadtler et al. changed the initial BCI mapping (blue) between neural activity and cursor movement to either remain inside (yellow) or move outside (red) that existing ten-dimensional neural manifold. When the altered activity patterns required to move the cursor retained trajectories within the existing neural manifold, monkeys could easily regain dexterous control within very few trials. However, when the changed mapping required the production of trajectories outside the existing manifold, monkeys could not easily regain control of the cursor within an experimental session.

This contrasting behaviour beautifully supported their hypothesis, providing the first causal hint that neural manifolds might not be mere mathematical descriptions of neural data: they may indeed capture something fundamental about how neural populations work. More than a decade later, mounting evidence across an increasing number of brain regions and species seems to support the view that, when interpreting brain activity, neural manifolds may be more than the sum of their neurons.

Juan Alvaro Gallego
Department of Bioengineering, Imperial College London, London, UK.
e-mail: jgallego@imperial.ac.uk

Competing interests

J.A.G. receives funding from Meta Platform Technologies, LLC and InBrain Neuroelectronics.

Original article: Sadtler, P. T. et al. Neural constraints on learning. *Nature* **512**, 423–426 (2014)

Related article: Stopfer, M. et al. Intensity versus identity coding in an olfactory system. *Neuron* **39**, 991–1004 (2003)