

Null ain't dull: new perspectives on motor cortex

Ta-Chu Kao and Guillaume Hennequin

Correspondence: g.hennequin@eng.cam.ac.uk (G. Hennequin)

Computational and Biological Learning lab, Department of Engineering,
University of Cambridge, Cambridge, UK

Abstract

Classical work has viewed primary motor cortex (M1) as a controller of muscle and body dynamics. A recent brain-computer interface experiment suggests a new, complementary perspective: M1 is itself a dynamical system under active control of other circuits.

Even the simplest of behaviours require concerted interactions among thousands of neurons. Of these many neurons, however, only a fraction directly determine behavioural outputs. For example, reaching for a cup of coffee can potentially be achieved by a myriad of different activity patterns in primary motor cortex (M1): as long as corticospinal (or “output-potent”) neurons produce the correct activity, the activity of other (“output-null”) neurons appears entirely unconstrained, or “redundant”. Redundancy has attracted much attention lately due to its potential significance for robust and flexible neural computations. Redundant representations improve robustness to perturbations [1, 2], might allow multiple computations to occur concurrently in the same circuit [3, 4], and could explain why behaviour remains stable despite routine reorganization of neural representations [5].

Importantly, neural redundancy could also hold important information concerning the circuit implementation of motor control. Indeed, although output-null activity does not *directly* contribute to behaviour, it is likely an essential cog in the mechanism that produces correct output-potent activity. Recently, Hennig et al. used a brain-computer interface (BCI) as a scientific tool to uncover the principles by which the brain chooses one pattern of output-null activity over another [6]. In monkeys, they recorded the activity of ~ 100 M1 neurons, used it as a control signal to actuate a cursor moving on a screen,

and trained the animals to perform specific cursor movements. Critically, this BCI setup allowed the authors to choose which linear combinations of the recorded neurons’ action potentials mattered for the cursor velocity, and which did not. In other words, they could arbitrarily create “output-potent” and “output-null” directions in the state space of neural activity, as illustrated in Figure 1A.

Hennig et al. used activity recorded during the BCI task to systematically rule out and rule in hypotheses regarding the structure of output-null activity in M1. A first possibility is that there is no predictable structure: M1 might receive noisy or task-unrelated inputs from other brain areas, and leave uncorrected these inputs’ contributions to output-null activity. Hennig et al. confronted two variants of this hypothesis to their data, and found that neither accurately predicted the distributions of activity along the output-null directions, across various directions of cursor movement.

A second hypothesis is inspired by previous work in motor neuroscience, in which M1 is typically viewed as controlling the dynamics of skeletal muscles (the “plant”) using appropriate inputs (Figure 1B; [7, 8]). According to well-established engineering wisdom, control inputs should ideally be kept small (relative to some nominal value) to ensure robustness of the control solution.

Strictly speaking, this principle applies to potent activity only (input to the muscular system). However, M1 might be implementing this principle more liberally and constrain its activity to be as “small” overall as the generation of correct potent activity permits. Hennig et al. considered two versions of this “minimal firing” hypothesis; again, neither made accurate predictions.

Substantially better predictions of output-null activity were obtained based on a third hypothesis seemingly unrelated to previous work in motor control. This “fixed distribution hypothesis” postulates that M1 tends to produce patterns of activity belonging to a fixed repertoire, which does not depend on the specific choice of potent directions. Given a choice of potent directions, activity is selected on a moment-by-moment basis from this fixed repertoire, on the condition that it elicit the right cursor velocity. Mathematically, this corresponds to conditioning a fixed distribution of M1 activity on some desired value of momentary potent activity (Figure 1D). Thus, if one knew the fixed distribution, one could predict the structure of output-null activity for any choice of potent/null directions. To test this hypothesis, Hennig et al. used activity recorded for one set of potent directions as an empirical proxy for the (unknown) fixed distribution, and used it to predict output-null activity under a second choice of potent directions. Remarkably, these predictions were better than those of any other hypothesis considered, and were as good as finite samples would allow.

While Hennig et al.’s fixed distribution hypothesis provides a compact, thought-provoking description of M1 activity, it lacks a computational rationale. What normative principle would account for the author’s observations, and illuminate the role of M1 in motor control? Hints might be found in recent experimental [9] and theoretical [10] work, in which M1’s complex activity patterns are understood as resulting from strong internal dynamics. Accordingly, beyond thinking of M1 as con-

trolling muscles (Figure 1B), one can view M1 as being part of the “plant”, i.e. an extension of the muscles that too needs to be controlled (presumably via control inputs from other brain areas; Figure 1C). Under this new perspective, the fixed distribution hypothesis emerges naturally. We illustrate this using a canonical model of cortical dynamics, with two coupled populations of excitatory and inhibitory cells (Figure 1E). Both populations receive inputs optimized for the production of some desired activity fluctuations along a chosen potent direction. From a control theoretic standpoint, strong network interactions imply that control inputs of fixed energy can steer activity further along some “preferred directions” than along others, by exploiting the network’s tendency to produce correlated activity patterns. Thus, if control inputs to M1 are energy-limited (as they are in our example, and as robustness demands), M1 activity under an optimal control policy should remain confined to a certain repertoire, or “fixed distribution” (Figure 1E, black ellipse). Importantly, this repertoire is a reflection of the network’s dynamics, and does not depend on the specific choice of potent directions. As expected, therefore, the fixed distribution hypothesis accurately predicts the statistical structure of output-null activity in this toy example. In other words, Hennig et al.’s findings are consistent with optimal control of M1 dynamics under energy constraints.

Going forward, we speculate that much will be learned about the neural basis of movement by thinking of M1 (and spinal cord circuits) not only as body controllers, but also as dynamical systems under the control of other neural circuits. This new perspective will suggest principled ways of elucidating the role of motor areas upstream of M1 (e.g. thalamic nuclei, basal ganglia, cerebellum). Examining neural redundancy at each level of the control hierarchy, e.g. using BCI-inspired techniques, will continue to bring useful insights: null ain’t dull under the skull.




Figure 1. Neural redundancy in M1 suggests a new view of M1 as a controlled dynamical system. (A) Illustration of neural redundancy: the same behaviour (natural, or BCI-driven) could be produced by very different trajectories in the state space of neural activity (three shown here). The activity along “potent” directions is constrained by the desired behaviour, and is therefore the same for all candidate trajectories (top right inset). In contrast, activity along “null” directions has no direct effect on behaviour and is therefore free to vary (top left inset). (B) M1-as-a-controller view. (C) M1-as-a-plant view. (D) Illustration of the fixed-distribution hypothesis. At any time, output-null activity is selected as if drawn from some fixed distribution of neural activity (heat map), conditioned on a momentary desired value of potent activity (white dot). (E) A 2-unit neural network (top) is driven by optimal control inputs to generate some desired fluctuations along a given potent direction. The distribution of network activity (dots) has the same structure irrespective of the potent direction being used (compare orange and green). The black ellipse delineates the region of state space within which the network activity can be steered given a fixed input energy budget.

References

- [1] Martin Boerlin, Christian K Machens, and Sophie Denève. Predictive coding of dynamical variables in balanced spiking networks. *PLoS Computational Biology*, 9:e1003258, 2013.
- [2] Nuo Li, Kayvon Daie, Karel Svoboda, and Shaul Druckmann. Robust neuronal dynamics in premotor cortex during motor planning. *Nature*, 532:459, 2016.
- [3] Matthew T Kaufman, Mark M Churchland, Stephen I Ryu, and Krishna V Shenoy. Cortical activity in the null space: permitting preparation without movement. *Nature Neuroscience*, 17:440, 2014.
- [4] Sergey D Stavisky, Jonathan C Kao, Stephen I Ryu, and Krishna V Shenoy. Motor cortical visuomotor feedback activity is initially isolated from downstream targets in output-null neural state space dimensions. *Neuron*, 95:195–208, 2017.
- [5] Laura N Driscoll, Noah L Pettit, Matthias Minderer, Selmaan N Chettih, and Christopher D Harvey. Dynamic reorganization of neuronal activity patterns in parietal cortex. *Cell*, 170:986–999, 2017.
- [6] Jay A Hennig, Matthew D Golub, Peter J Lund, Patrick T Sadtler, Emily R Oby, Kristin M Quick, Stephen I Ryu, Elizabeth C Tyler-Kabara, Aaron P Batista, Byron M Yu, and Steven M Chase. Constraints on neural redundancy. *eLife*, 7:e36774, 2018.
- [7] Emanuel Todorov. Direct cortical control of muscle activation in voluntary arm movements: a model. *Nature neuroscience*, 3:391, 2000.
- [8] Timothy P Lillicrap and Stephen H Scott. Preference distributions of primary motor cortex neurons reflect control solutions optimized for limb biomechanics. *Neuron*, 77:168–179, 2013.
- [9] Krishna V Shenoy, Maneesh Sahani, and Mark M Churchland. Cortical control of arm movements: a dynamical systems perspective. *Annual review of neuroscience*, 36, 2013.
- [10] Guillaume Hennequin, Tim P Vogels, and Wulfram Gerstner. Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron*, 82:1394–1406, 2014.