

Spotlight

Null Ain't Dull: New Perspectives on Motor Cortex

Ta-Chu Kao¹ and
Guillaume Hennequin^{1,*}

Classical work has viewed primary motor cortex (M1) as a controller of muscle and body dynamics. A recent brain-computer interface (BCI) 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. However, of these many neurons, only a fraction directly determine behavioural outputs. For example, reaching for a cup of coffee can potentially be achieved by myriad 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 reorganisation 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 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, the authors 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 action potentials of the recorded neurons 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.

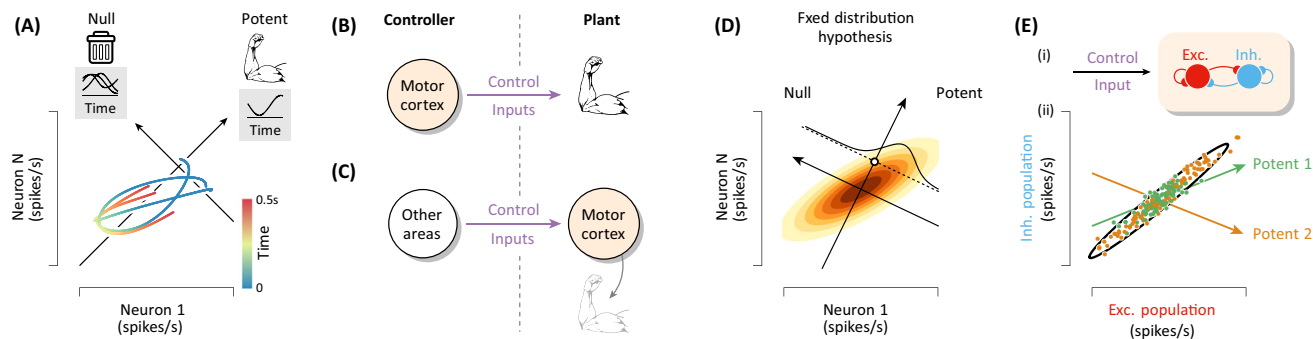
Hennig *et al.* used activity recorded during the BCI task to systematically rule out and rule in hypotheses regarding the structure of output-null activity in M1 [6]. A first possibility is that there is no predictable structure: M1 might receive noisy or task-unrelated inputs from other brain areas, and leave uncorrected the contributions of these inputs to output-null activity. Hennig *et al.* tested two variants of this hypothesis with 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 [6].

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 elicits 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 and/or 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 their observations, and illuminate the role of M1 in motor control? Hints might be found in recent experimental [9] and theoretical [10] work, in which the complex activity patterns of M1 are understood as resulting from strong internal dynamics. Accordingly, beyond thinking of M1 as controlling muscles (Figure 1B), one can view M1 as being part of the 'plant' [i.e., an extension of the muscles



Trends in Cognitive Sciences

Figure 1. Neural Redundancy in Primary Motor Cortex (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 different trajectories in the state space of neural activity (three shown here). The activity along ‘potent’ directions is constrained by the desired behaviour and, therefore, is the same for all candidate trajectories (top-right inset). By contrast, activity along ‘null’ directions has no direct effect on behaviour and, therefore, is 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 though drawn from some fixed distribution of neural activity (heat map), conditioned on a momentary desired value of potent activity (white dot). (E) A two-unit neural network (i) is driven by optimal control inputs to generate some desired fluctuations along a given potent direction. (ii) 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.

that also 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 optimised 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 tendency of the network 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 dynamics of the network, and does not depend on the specific choice of potent directions. Therefore, as

expected, 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 [6] 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 a body controller, but also as a dynamical system 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, and 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.

Supplemental Information

Supplemental information associated with this article can be found, in the online version, at <https://doi.org/10.1016/j.tics.2018.09.005>.

¹Computational and Biological Learning Lab, Department of Engineering, University of Cambridge, Cambridge, UK

*Correspondence:
g.hennequin@eng.cam.ac.uk (G. Hennequin).
<https://doi.org/10.1016/j.tics.2018.09.005>

References

- Boerlin, M. *et al.* (2013) Predictive coding of dynamical variables in balanced spiking networks. *PLoS Comput. Biol.* 9, e1003258
- Li, N. *et al.* (2016) Robust neuronal dynamics in premotor cortex during motor planning. *Nature* 532, 459
- Kaufman, M.T. *et al.* (2014) Cortical activity in the null space: permitting preparation without movement. *Nat. Neurosci.* 17, 440
- Stavisky, S.D. *et al.* (2017) Motor cortical visuomotor feedback activity is initially isolated from downstream targets in output-null neural state space dimensions. *Neuron* 95, 195–208
- Driscoll, L.N. *et al.* (2017) Dynamic reorganization of neuronal activity patterns in parietal cortex. *Cell* 170, 986–999
- Hennig, J.A. *et al.* (2018) Constraints on neural redundancy. *eLife* 7, e36774
- Todorov, E. (2000) Direct cortical control of muscle activation in voluntary arm movements: a model. *Nat. Neurosci.* 3, 391
- Lillcrap, T.P. and Scott, S.H. (2013) Preference distributions of primary motor cortex neurons reflect control solutions optimized for limb biomechanics. *Neuron* 77, 168–179
- Shenoy, K.V. *et al.* (2013) Cortical control of arm movements: a dynamical systems perspective. *Annu. Rev. Neurosci.* 36, 337–359
- Hennequin, G. *et al.* (2014) Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron* 82, 1394–1406