Crouching tiger, hidden dimensions

Terence D Sanger & John F Kalaska

A study finds that, during movement preparation, when motor cortex is active, but elicits no muscle output, firing of individual neurons in dorsal premotor and primary motor cortex cancels out at the level of population activity.

We have all seen images of a cat stalking its prey, crouched, motionless but ready to spring into action. How does an animal that is preparing to make a movement hold still until it is time to move? An elegant study of neural activity in the dorsal premotor (PMd) and primary motor cortex (M1) of nonhuman primates by Kaufman *et al.*[1] provides an intriguing and provocative new hypothesis about the cerebral cortical neural mechanisms underlying the transition from a state of motor preparation to overt action: neural activity patterns during preparation are structured so that they cannot cause muscle contractions—they are 'output null'—whereas the neural activity that initiates movement are 'output potent'—they can cause muscle contractions when summed at the spinal cord level.

The transition from a stable posture to externally visible action is initiated and controlled by massively distributed patterns of neuronal activity in many cortical and subcortical structures, culminating in the activation of spinal motor neurons and muscle contractions. However, those same movement-related neurons can also be very active even during periods of physical inactivity. One can imagine or even intend to make an arm movement while maintaining the limb in a fixed posture. Covertly imagined movements and overtly performed movements activate the same brain regions[2]. Neural activity that encodes different attributes of an intended movement is generated in many movement-related structures during the delay period between the presentation of a cue that instructs a subject to prepare to make that movement and the arrival of a second cue that tells the subject to initiate the movement, but does not cause any measureable change in muscle activity[3,4]. Similarly, cerebral cortical 'mirror' neurons are active not only when subjects make particular movements, but also when they observe others perform the same actions[5]. Mirror-like activity has even been observed in M1 corticospinal neurons[6]. Why doesn't all of this activity in motor-related neural structures lead to real muscle activation?

The answer in sleep is known: there are powerful brainstem inhibitory mechanisms that block the imagined movements of dreams from activating spinal cord motor circuits[7]. These are the mechanisms that partially fail in people who sleepwalk and that may be overactive in people with sleep paralysis[7]. But such mechanisms are all or none: you can either move anything or you move nothing. Lack of movement in these cases includes a flaccid lack of posture, as muscles are relaxed and are unable to actively resist perturbation any more than they can actively produce movement. Such a mechanism seems unlikely to explain lack of motor output during conscious imagined movements or during preparation for upcoming movement. It is not known whether the skeletomotor system has a brake mechanism similar to the brainstem omnipause neurons that control when voluntary eye saccades are launched[8]. Attempts to identify a corresponding brake mechanism in PMd and M1 have been largely unsuccessful[9,10]. However, descending corticospinal signals onto spinal interneurons could potentially serve to prevent premature activation of spinal motor neurons during supraspinal preparatory activity[6,11].

Kaufman *et al.*[1] propose a different mechanism (**Fig. 1**). Their hypothesis begins with a few simple, but reasonable, assumptions. First, muscle activity patterns are generated by an approximately linear weighted summation in the spinal cord of the descending supraspinal spike trains from many cortical motor neurons. Second, the presence or absence of muscle activity and overt movement depends solely on that process of linear summation of descending motor signals at the spinal level; there is no need to invoke a parallel and independent brake mechanism (**Fig. 1**).

Because there are many more cortical motor neurons than muscles, there must be many different patterns of motor cortical neural firing that can produce the same muscle activations. Kaufman *et al.*[1] called these particular cortical neural activity patterns 'output potent' to signify that the generation of any one of these patterns in the cortex will result in muscle contractions and movement when summed in the spinal cord. Similarly, those same populations of motor cortical neurons have many options for complex time-varying firing patterns that do not cause changes in muscle activation when summed in the spinal cord. These 'output-null' response patterns can be used to prime certain neural circuits during motor preparation, but will not cause overt changes in muscle contractile activity because the patterns of neural activity, by design, cancel each other out when they converge on spinal motor neurons. Because an output-null pattern can be added to an ongoing pattern of background activity without affecting muscle activation, this provides a mechanism for preventing movement while maintaining posture and preparing to move. To initiate the movement, the activity of cortical populations must change to output-potent patterns at the desired moment of movement onset.

The hypothesis predicts that cortical neural activity during motor preparation is functionally different than that during movement: we do not prepare to move by covert mental rehearsal of the neural pattern that causes the intended movement. Indeed, the hypothesis implies that it is not possible to avoid movement while generating a true copy of the movement-related neural activity during preparation; avoidance of movement is possible only by avoiding the output-potent patterns of neural activity that cause movement.

To test these predictions, Kaufman *et al.*[1] first used sophisticated numerical tools to identify the output-potent dimensions of neural activity in PMd and M1 of rhesus monkeys from neural activity recorded during an instructed-delay arm reaching task. They then showed that the neural activity during the preparatory period before movement has more of its energy in output-null dimensions orthogonal to (and thus uncorrelated with) these output-potent dimensions. It appears that the brain knows which patterns cause movement and is able to avoid those patterns when needed while still generating neural activity patterns that prepare the movement.

Although these findings are largely consistent with the hypothesis, they also suggest that

Terence D. Sanger is in the Departments of Biomedical Engineering, Biokinesiology and Child Neurology, University of Southern California, Los Angeles, California, USA, and John F. Kalaska is in the Département de Neurosciences, Faculté de médecine, Université de Montréal, Montréal, Québec, Canada.
e-mail: terry@sangerlab.net or john.francis.kalaska@umontreal.ca

it may not be the entire story. The output-null dimensions accounted for most of the activity generated during the preparatory period. However, the remainder was still expressed in the output-potent dimensions, but no measurable change in muscle activity was observed. Although this might be a result of simplifying assumptions or other shortcomings in their model, it may also indicate that other processes, including a brake mechanism, might still be involved.

This begs the question of what motor preparation actually does, as the hypothesis requires that the brain cannot use the same neural activity patterns that are used during movement. That is no doubt a prime topic for future work. Previous studies by this same group[12] and many others have provided some insight into the nature of cortical motor preparatory activity. If we believe that the brain does not 'waste' neurons and neural activity, then it seems unlikely that output-null dimensions truly represent patterns that accomplish nothing; these are simply patterns that do not directly contribute to muscle activation during the particular experimental task.

The existence of output-null dimensions depends on the large disparity between the

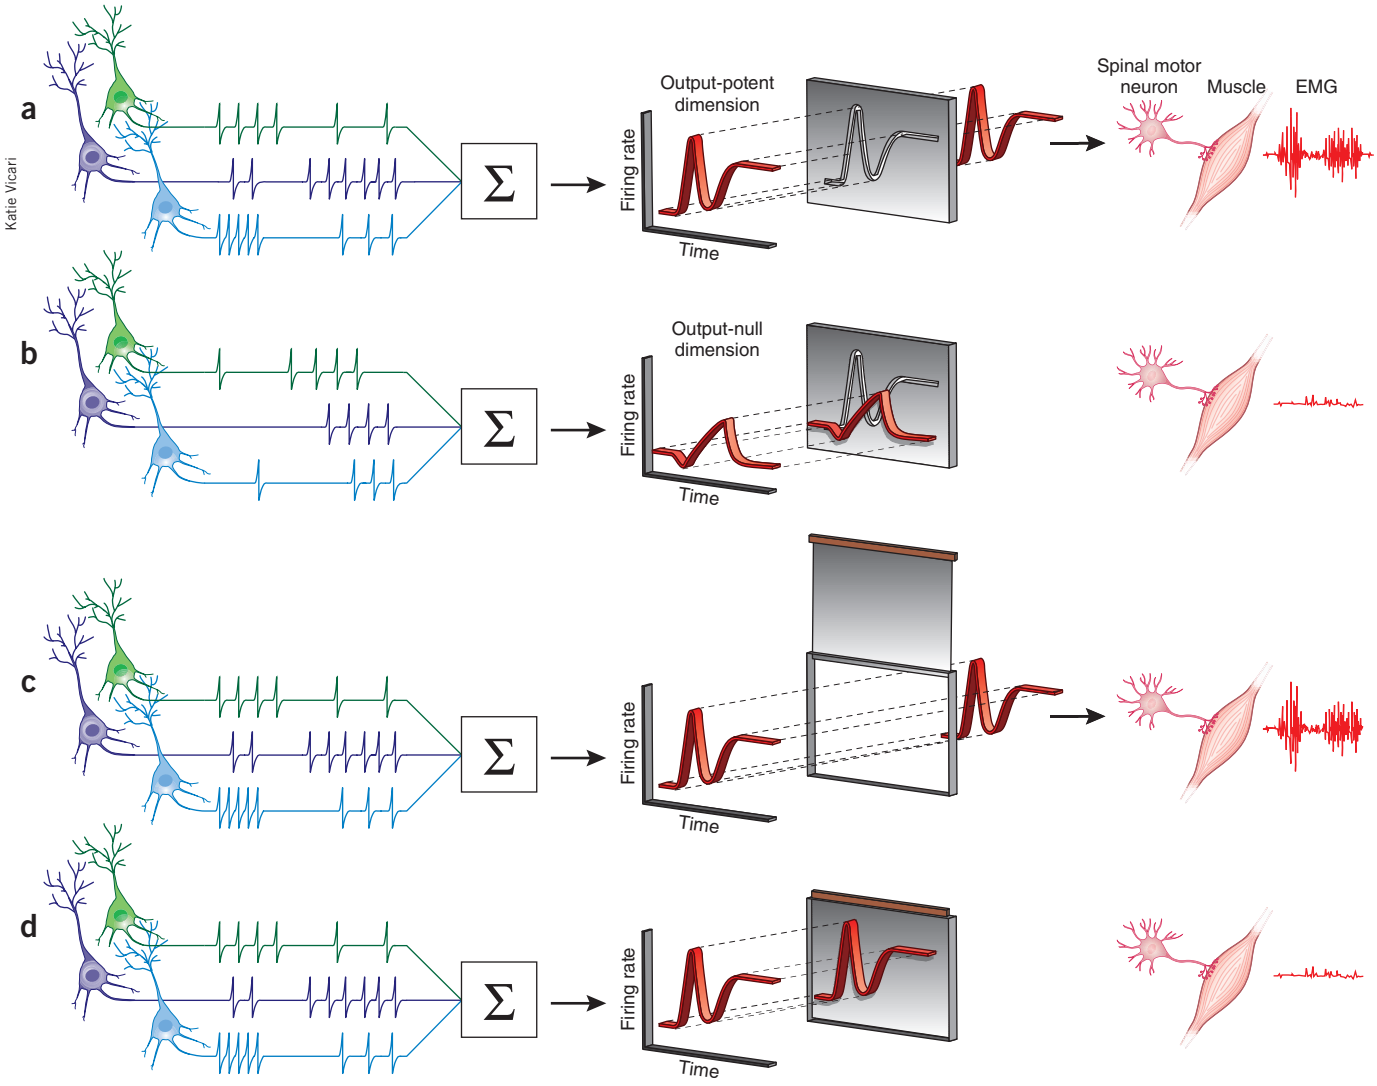

Figure 1 A highly schematic illustration of the difference between the output-potent versus output-null hypothesis and the motor gating hypothesis. (**a**,**b**) In the potent versus null hypothesis, only certain patterns of motor cortical activity affect muscles (**a**); different patterns of descending cortical activity cancel out and have no effect (**b**). In this hypothesis, whether or not motor cortical activity affects muscles is determined by the pattern of the neural activity itself. (**c**,**d**) In the motor gating hypothesis, in contrast, signals are transmitted to the muscles only when a gate is open (**c**) and are otherwise blocked even if the pattern of cortical activity would normally cause a muscle contraction (**d**). In this hypothesis, whether or not neural activity affects muscles is determined by the gate and does not depend on the pattern of neural activity. This illustration is intended only to contrast a particular feature of the two hypotheses; the actual mechanism of identifying output-null versus output-potent activity patterns proposed by the authors is determined by the mathematical projection of descending cortical activity onto a high-dimensional subspace and should not be construed as necessarily matching a particular temporal pattern. The plane in **a** and **b** represents the subspace of all possible cortical discharge patterns (dimensions). When a discharge pattern occupies an output-potent dimension, illustrated graphically as a slot in the subspace plane, it can access spinal motor neurons and cause a muscle contraction (**a**). Other patterns cannot pass through (**b**). In essence, the high-dimensional subspace acts like a filter to allow some cortical activity to cause muscle contractions, whereas other activity patterns cannot. The potent versus null hypothesis does not stipulate where that process occurs. For simplicity, the authors describe spinal motor neurons as the site at which descending cortical signals converge, but it is just as likely that spinal interneuronal networks and other neural circuits also contribute to the filtering of output-potent versus output-null discharge patterns.

dimensionality of the neural space and the dimensionality of the muscle space. Such disparities lead to redundancy, with different neural patterns causing the same muscle activity while others do not. Kaufman *et al.*[1] propose that the redundant neural dimensions are tightly controlled for the specific purpose of motor preparation. In a contrasting example of the exploitation of redundancy in the motor system, the 'uncontrolled manifold' hypothesis suggests that redundant biomechanical dimensions are not controlled and are allowed to vary almost randomly if they do not affect performance[13].

One of the striking findings of the analysis is that the entire theory is consistent with a linear or near-linear model relating cortical neural population activity to motor control. The anatomy and biomechanical properties of the primate skeletomotor system are notoriously complex and nonlinear. Any motor controller designed on first principles derived from physics that would attempt to match the performance of the biological motor system would have to be equally complex and nonlinear. Nevertheless, as with all studies of complicated nonlinear systems, it is often best to start by determining how far we can go with a much simpler linear model. It is surprising that the output-potent and output-null dimensions can be extracted by a linear regression between neural activity and muscle activity and that they are truly orthogonal in the sense predicted by linear algebra. This did not have to be the case: the two dimensions could have been twisted through the space of possible neural activations in intricate ways. And yet the output-null dimension can be approximately described by a hyperplane that is the null space of a linear (matrix) transformation between cortical activity and muscle activity. This is reminiscent of the success of the population-vector analysis, for which the relation between motor cortical firing and movement kinematics is approximately linear[14]. This near linearity might indicate that cortical control processes are a "good enough" solution[15] that provides subcortical motor circuits with enough information to initiate an approximation of the desired action, whereas the latter use feedforward and feedback signals to deal with motor errors and the nonlinearities of the peripheral skeletomotor apparatus. Moreover, this seeming simplicity might reflect the influence of evolutionary processes that acted to optimize both the peripheral motor apparatus and its central neural controller as they coevolved over millions of years.

The mechanism proposed by Kaufman *et al.*[1] is necessarily a property of populations of neurons, rather than individual neurons, and it represents a very interesting example of the emerging role of theories that are able to predict behavior from simultaneously recorded neural populations rather than from single cells. The existence and role of the output-null dimensions would not have been detectable from single-cell recordings. Indeed, the authors point out that the same cell may participate in both the output-potent and output-null dimensions at different times, even though in the latter case its potential effect on muscles must be cancelled out by the activity of other neurons in the population. As multielectrode recordings become more and more prevalent, the availability and importance of population-based theories of neural encoding will become increasingly evident. Just as multiple visual images can be overlaid and perceived simultaneously, so populations of neurons have tremendous power to represent more than one perception or movement simultaneously. Perhaps this is how we can sit still without acting out our every motor whim whenever we imagine doing something, and how the hunting tiger can crouch motionless while preparing to pounce.

COMPETING FINANCIAL INTERESTS
The authors declare no competing financial interests.

1. Kaufman, M.T., Churchland, M.M., Ryu, S.I. & Shenoy, K.V. *Nat. Neurosci.* **17**, 440–448 (2014).
2. Hétu, S. *et al. Neurosci. Biobehav. Rev.* **37**, 930–949 (2013).
3. Bremner, L.R. & Andersen, R.A. *Neuron* **75**, 342–351 (2012)
4. Buch, E.R., Brasted, P.J. & Wise, S.P. *Exp. Brain Res.* **169**, 69–84 (2006).
5. Rizzolatti, G. & Sinigaglia, C. *Nat. Rev. Neurosci.* **11**, 264–274 (2010).
6. Vigneswaran, G., Philipp, R., Lemon, R.N. & Kraskov, A. *Curr. Biol.* **23**, 236–243 (2013).
7. Hishikawa, Y. & Shimizu, T. *Adv. Neurol.* **67**, 245–271 (1995).
8. Evinger, C., Kanako, C.R. & Fuchs, A.F. *J. Neurophysiol.* **47**, 827–844 (1982).
9. Kaufman, M.T. *et al. J. Neurophysiol.* **104**, 799–810 (2010).
10. Kaufman, M.T., Churchland, M.M. & Shenoy, K.V. *J. Neurophysiol.* **110**, 817–825 (2013).
11. Cohen, O., Sherman, E., Zinger, N., Perlmutter, S. & Prut, Y. *Curr. Opin. Neurobiol.* **20**, 696–703 (2010).
12. Afshar, A. *et al. Neuron* **71**, 555–564 (2011).
13. Scholz, J.P. & Schöner, G. *Exp. Brain Res.* **126**, 289–306 (1999).
14. Georgopoulos, A.P., Kettner, R.E. & Schwartz, A.B. *J. Neurosci.* **8**, 2928–2937 (2008).
15. Partridge, L.D. *Am. J. Physiol.* **242**, R172–R177 (1982).

Oxytocin for all senses

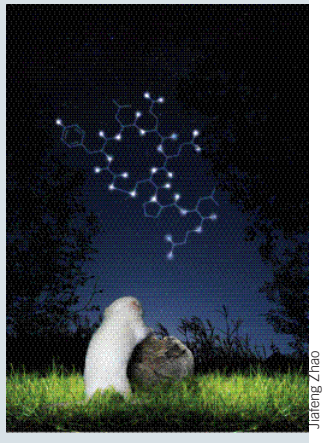
Jialeng Zhao

Deprivation of one sensory modality is known to diminish cortical responses in the corresponding sensory cortex as well as alter neuronal responses in unrelated sensory cortices. Yet little is known about what drives this cross-modal plasticity. On page 391, Zheng and colleagues explore the effects of early unimodal sensory deprivation in multiple sensory cortical regions in mice and report that oxytocin, a neuropeptide important for lactation, parturition, social and emotional behaviors, is critical for cross-modal cortical plasticity.

Zheng and colleagues subjected mice to sensory deprivation starting from birth by either removing their whiskers or raising them in the dark. Each type of deprivation decreased spontaneous excitatory synaptic transmission and sensory stimulation-evoked responses in both primary somatosensory (S1) and primary visual (V1) cortex but not in higher order cortical regions such as the prefrontal cortex. This was accompanied by a decrease in oxytocin levels in the sensory cortex affected by the sensory deprivation. Whisker removal or dark rearing also reduced oxytocin production and expression in the hypothalamus. Oxytocin injection into S1 not only increased excitatory synaptic transmission in both S1 and V1 but also reversed the cross-modal effects of sensory deprivation in these areas. To complement their findings from sensory deprivation, Zheng *et al.* then reared mice in a sensory-enriched environment. This manipulation led to increased oxytocin production in the hypothalamus, higher levels of oxytocin expression in S1 and V1, and enhanced neuronal responses in these cortical areas. Sensory enrichment also rescued the effects of sensory deprivation, similar to the treatment with exogenous oxytocin.

These findings reveal a critical function for oxytocin in activity-dependent cortical development and cortical plasticity.

Jean Mary Zarate