Box 2. Who Is Allowed to Experience Outrage?

We have briefly discussed how pro-empathy rhetoric can be leveraged to delegitimize outrage, particularly among marginalized groups. Indeed, promoting intergroup harmony can reinforce an inequitable social structure: if conditions appear harmonious, high-status groups feel reassured that the status quo is fair and low-status groups feel their grievances are less legitimate [13]. This phenomenon is compounded by the observation that only certain groups are ‘allowed’ to express outrage. For example, stigmatized group members are often held to higher moral standards (e.g., accused of expressing inappropriate emotions, especially anger, at greater rates than majority group members [14]). In short, people often put boundary conditions on who is permitted to experience outrage. The challenge is that outrage is only effective for promoting collective action if people are allowed—and allow themselves—to feel it.

Power, status, and majority/minority group membership can all determine who is likely to express outrage versus suppress it, but this relationship may be bidirectional. In other words, expression of outrage may not just motivate collective action, but also help define the collective itself through the creation of common cause. Sharing outrage with others may act as a group-level emotion, facilitating further cohesion [15] and amplifying outrage’s potency for motivating behavior.

become imbued with positive versus negative connotations? Again, we believe the most fruitful approach to answering these questions involves bringing moral psychology and intergroup researchers together.

Concluding Remarks

In both public discourse and in psychological research, outrage is frequently cast in a destructive light. In contrast, the intergroup literature presents several cases in which outrage can serve as an important catalyst for collective action. There is no shortage of current events that demonstrate how effective outrage can be at uniting people in democracy-preserving behavior, but we need a better account of its dynamics and efficacy in light of its costs. In merging the intergroup and moral psychology literatures, we hope to promote a more complete view of outrage—as an emotion that might lead to interpersonal antagonism, but that may also act as a lever for activism on a societal scale.

Acknowledgments

The writing of this paper was facilitated by a National Science Foundation Graduate Research Fellowship awarded to V.L.S., NSF 1660707 awarded to C.D.C., and NSF 1551559 awarded to M.C.

Resources

[i] https://www.history.com/this-day-in-history/second-moratorium-against-the-war-held

[ii] https://www.theatlantic.com/technology/archive/2017/01/womens-march-protest-count/514166/

[iii] https://nypost.com/2018/09/04/nike-shares-drop-amid-backlash-over-new-kaepernick-ad/

[1] Department of Psychology, The Pennsylvania State University, University Park, PA, USA
[2] Rock Ethics Institute, The Pennsylvania State University, University Park, PA, USA
[3] Department of Psychology, Harvard University, Cambridge, MA, USA

*Correspondence: vls23@psu.edu (V.L. Spring).

https://doi.org/10.1016/j.tics.2018.09.006

References

1. Zaki, J. (2017) Moving beyond stereotypes of empathy. *Trends Cogn. Sci.* 21, 59–60
2. Gummerum, M. *et al.* (2016) Costly third-party interventions: the role of incidental anger and attention focus in punishment of the perpetrator and compensation of the victim. *J. Exp. Soc. Psychol.* 65, 94–104
3. Crockett, M.J. (2017) Moral outrage in the digital age. *Nat. Hum. Behav.* 1, 769–771
4. Rothschild, Z.K. and Keefer, L.A. (2017) A cleansing fire: moral outrage alleviates guilt and buffers threats to one’s moral identity. *Motiv. Emot.* 41, 209–229
5. van Kleef, G.A. and Fischer, A.H. (2016) Emotional collectives: how groups shape emotions and emotions shape groups. *Cogn. Emot.* 30, 3–19
6. Tagar, M. *et al.* (2011) The positive effect of negative emotions in protracted conflict: the case of anger. *J. Exp. Soc. Psychol.* 47, 157–164
7. Becker, J.C. and Wright, S.C. (2011) Yet another dark side of chivalry: benevolent sexism undermines and hostile sexism motivates collective action for social change. *J. Pers. Soc. Psychol.* 101, 62–77
8. Ford, B.Q. *et al.* (2018) Using reappraisal to regulate negative emotion after the 2016 US presidential election: does emotion regulation trump political action? *J. Pers. Soc. Psychol.* Published online June 28, 2018. http://dx.doi.org/10.1037/pspp0000200
9. Sawaoka, T. and Monin, B. (2018) The paradox of viral outrage. *Psychol. Sci.* Published online August 9, 2018. http://dx.doi.org/10.1177/0956797618780658
10. Netzer, L. *et al.* (2015) Interpersonal instrumental emotion regulation. *J. Exp. Soc. Psychol.* 58, 124–135
11. Greene, J. (2014) *Moral Tribes: Emotion, Reason, and the Gap Between Us and Them,* Penguin
12. Shteynberg, G. (2017) A collective perspective: shared attention and the mind. *Curr. Opin. Psychol.* 23, 93–97
13. Dixon, J. *et al.* (2016) From prejudice reduction to collective action: two psychological models of social change (and how to reconcile them). In *Cambridge Handbook of the Psychology of Prejudice* (Barlow, F. and Sibley, C., eds), pp. 481–499, Cambridge University Press
14. Jones, T. and Norwood, K.J. (2016) Aggressive encounters & White fragility: deconstructing the trope of the angry Black woman. *Iowa Law Rev.* 102, 2017–2069
15. Smith, E.R. and Mackie, D.M. (2015) Dynamics of group-based emotions: insights from intergroup emotions theory. *Emot. Rev.* 7, 349–354

Spotlight

Null Ain’t Dull: New Perspectives on Motor Cortex

Ta-Chu Kao[1] 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 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

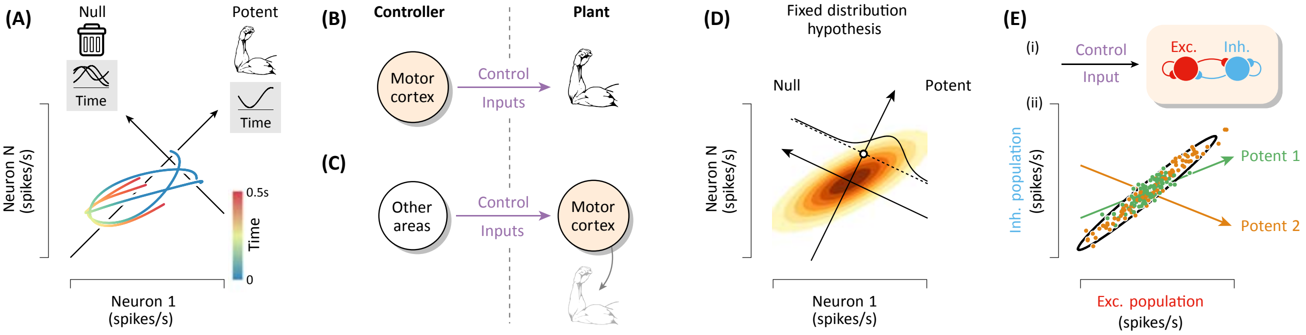

Trends in Cognitive Sciences

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.

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.

Acknowledgment

G.H. is supported by Wellcome Trust Seed Award n. 202111/Z/16/Z.

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.

[1]Computational 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

1. Boerlin, M. *et al.* (2013) Predictive coding of dynamical variables in balanced spiking networks. *PLoS Comput. Biol.* 9, e1003258
2. Li, N. *et al.* (2016) Robust neuronal dynamics in premotor cortex during motor planning. *Nature* 532, 459
3. Kaufman, M.T. *et al.* (2014) Cortical activity in the null space: permitting preparation without movement. *Nat. Neurosci.* 17, 440
4. 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
5. Driscoll, L.N. *et al.* (2017) Dynamic reorganization of neuronal activity patterns in parietal cortex. *Cell* 170, 986–999
6. Hennig, J.A. *et al.* (2018) Constraints on neural redundancy. *eLife* 7, e36774
7. Todorov, E. (2000) Direct cortical control of muscle activation in voluntary arm movements: a model. *Nat. Neurosci.* 3, 391
8. Lillicrap, 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
9. Shenoy, K.V. *et al.* (2013) Cortical control of arm movements: a dynamical systems perspective. *Annu. Rev. Neurosci.* 36, 337–359
10. Hennequin, G. *et al.* (2014) Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron* 82, 1394–1406

Spotlight

Extraordinary Altruism and Transcending the Self

Molly J. Crockett[1,*] and Patricia L. Lockwood[2]

Longstanding psychological theories posit a link between empathy and altruism. A new study of anonymous kidney donors finds these 'extraordinary altruists' show an increased overlap in neural responses to pain for self and others. These findings, alongside other recent studies of altruism, shed new light on the nature of selflessness.

On 6 June 2014, 21-year-old Michael Campbell was driving his normal delivery route when he witnessed a nearby house burst into flames. Hearing cries for help, Campbell rushed into the burning building and discovered a badly injured man covered in debris in the second-floor bedroom. Campbell dragged the man down the stairs and out of the house just as the second floor of the house collapsed. Both men survived and Campbell received the Carnegie Medal for his heroic actions. Extraordinary acts of selflessness like this are captivating and inspiring, and a deeper scientific understanding of what makes someone risk their life for a stranger has the potential to transform