NOTES

1. Those interested in more technical details should see any of the many works that discuss KFs in detail, for example, Kalman (1960); Kalman and Bucy (1961); Gelb (1974); Bryson and Ho (1969); and Haykin (2001); for some discussion of applications of KFs (among other constructs) to understanding brain function, see Eliasmith and Anderson (2003).

2. It might be wondered what justification there is for assuming that the driving force can be predicted accurately. This is just by definition. It is assumed that the process is subject to external influences. Any influence that is completely predictable is a driving force; the rest of the external influence – whatever is not predictable – is process noise. So in a case where there were an "unpredictable" driving force, this would actually be part of the process noise.

3. See discussion in Hutchins (1995; especially Ch. 3).

4. What gets suppressed is the overt performance. Interestingly, however, a great many other bodily events normally associated with overt performance, such as increases in metabolic activity, heart rate, and so on, accompany many kinds of motor imagery. For a review, see Jeannerod (1994). In this article, when I speak of motor commands being suppressed in favor of the processing of an efference copy, I mean only the overt bodily movements are suppressed. It may even be so that in some cases there is a small degree of muscular excitation, perhaps because the motor signals are not completely blocked.

5. It should be noted that Johnson's position here is not exactly the same as Jeannerod's, because he claims that this imagery is used in order to construct a final motor plan. But, as far as I can tell anyway, Johnson nevertheless is maintaining that it is imagery that is being used, and that this imagery is the result of the "simulated" operation of efferent motor areas, those involved in planning a movement. The details are complex, though, and Johnson's position may not be a good example of what I call the simulation theory.

6. The situation here is complex. It is not clear to what extent and under what conditions the MSS emulator adapts as a function of plant drift. While the case of phantom limb patients suggests that it can, other cases of paralysis suggest that this is not always so (Johnson 2000b). I will simply note that the emulation theory itself need not take a stand on whether, and under what conditions, emulators are malleable. I use the example of apparent malleability in the case of phantom limb patients to make the contrast between the emulation and the simulation theories clear. But that clarificatory role does not depend on the empirical issue of the conditions under which such malleability actually obtains.

7. A benefit of the emulation theory over the simulation theory is that it allows us to make sense of the difference between (a) things which we cannot move but *do not feel paralyzed* and (b) things which we cannot move and do *for that reason* feel paralyzed. The first group includes not only our own body parts over which we have no voluntary control, such as our hair, but also foreign things such as other people's arms, chairs and tables, et cetera. We cannot move these things, but the phenomenology of their not being voluntarily movable is not like that of a paralyzed part. If mere lack of ability to produce a motor plan accounted for the feeling of paralysis, then all of these things should seem paralyzed. On the emulation theory, the feeling of paralysis is the product of a mismatch between a motor plan and the resultant feedback, whether from the body or the emulator. Such a mismatch is possible only when we can produce a motor plan that mismatches the result of the attempt to effect that motor plan.

8. One difference is that Murphy was an actual robot, whereas the model discussed here is completely virtual.

9. Mel's Murphy is a very simple system working in a very constrained environment. More complex environments, including those with objects that moved without the agent willing it, would be far less predictable. This is of course much of the reason why, in perception, the Kalman gain is set fairly high. I use Murphy because its simplicity makes it a good exemplar for introducing the basic ideas, and I simply note that real perceptual situations will require much more sophistication. For an example of the sort of complexities that a full version of this sort of mechanism would need to deal with, see Nolfi and Tani (1999) and the references therein.

10. The emulation theory predicts exactly Wexler et al.'s results. Thus, if it were the case that motor areas were not active during "active" visual imagery, or if it were the case that the specific nature of the motor command associated with the imagined movement (rotate right vs. rotate left, for example) were not recruited during such imagery, then this would be prima facie evidence against the emulation theory.

11. I call this kind of imagery *amodal* rather than *multimodal* because this sort of imagery, if it in fact exists, is not tied to any modality. But in a sense it is multimodal, because it can be used to produce a modal image in any modality so long as a measurement procedure appropriate to that modality is available. The expressions *amodal* and *multimodal* are used in many ways, and it may be that what I here am calling amodal imagery might be close to what some researchers have called multimodal imagery.

12. For example: If x is between a and b, and b is between a and c, is x necessarily between a and c? There is reason to think that such questions are answered by engaging in spatial imagery, but little reason to think that much in the way of specifically visual mock experience is needed, though of course it might be involved in specific cases.

13. What I have in mind here is the idea that the neurally implemented emulator represents states by things like firing frequencies, phases, and such. A neural pool that is representing the presence of a predator behind the rock by firing rapidly can be "directly measured" in the sense that other neural systems can be wired such as to sniff that pool's activation state, and hence be sensitive to the presence of the predator. A "measurement" of this state would yield a visual image of a rock, because the predator is not in the visual image, and hence the narrowly modal emulator would throw away relevant information.

14. Exactly how to understand such a system is not trivial. Understanding the *which* system as an attentional tagging mechanism is sufficient for present expositional purposes, but my suspicion, which I am not prepared to argue for here, is that it is a system that has richer representational properties, such as the constitution of basic object identity. Of course, the richer sort of mechanism, if there is one, will surely be based at least in part upon a simpler attentional tagging mechanism.

15. I owe this phrase to Ramesh Jain, who produced it during a talk at UCSD.

Open Peer Commentary

Redundancy in the nervous system: Where internal models collapse

Ramesh Balasubramaniam

Sensory Motor Neuroscience, School of Psychology, University of Birmingham, Edgbaston, B15 2TT, United Kingdom.

r.balasubramaniam@bham.ac.uk **http://www.bham.ac.uk/symon**

Abstract: Grush has proposed a fairly comprehensive version of the idea of internal models within the framework of the emulation theory of representation. However, the formulation suffers from assumptions that render such models biologically infeasible. Here I present some problems from physiological principles of human movement production to illustrate why. Some alternative views to emulation are presented.

In the target article, Grush presents a unified theory for psychology based on the idea that the nervous system uses internal mod-

https://doi.org/10.1017/S0140525X0422009X Published online by Cambridge University Press

els that emulate the input/output relationships between sensory signals, actions, and their mutual consequences. The emulation theory advanced by the author is based on two aspects: forward internal models and inverse dynamics. The notion of forward internal models, which has drawn from work in adaptive control, arises from the idea that the nervous system takes account of dynamics in motion planning. Inverse dynamics is a clever means to establish the joint torques necessary to produce desired movements. I will now illustrate the failure of emulation-based models when dealing with issues of redundancy: a fundamental problem that the nervous system faces in assembling the units of action.

Redundancy problems in movement organization. The number of available degrees of freedom (DFs) of the body is typically greater than that required to reach the motor goal (DFs redundancy). The number of muscles per one DF is much greater than two (multimuscle redundancy). The abundance of DFs almost always makes a variety of solutions available to the nervous system in any given situation. Thus, a motor goal may be achieved differently depending on our intentions, external environmental (e.g., obstacles), or intrinsic (neural) constraints. Despite this flexibility, the control of actions is unambiguous: Each time the body moves, a unique action is produced despite the possibility of using other actions leading to the same goal. It is unclear how these seemingly opposite aspects – flexibility and uniqueness – are combined in the control of actions. Following Bernstein (1967), we refer to these aspects of action production as the "redundancy problem."

Computational problems: Multi-joint redundancy and inverse solutions. One of the inherent assumptions about motor control in the emulation theory is the central specification of output variables (e.g., force or muscle activation patterns). Moreover, it is supposed that these output variables are made routinely available to the nervous system through a combination of inflow and outflow signals. I argue that internal models cannot deal with redundancies in the nervous system. In fact, internal models bring additional layers of redundancy to the system at each level of the nervous system. To demonstrate this I will present a simple example from the inverse dynamic computations for multi-joint arm movements. A fundamental assumption made here is that the computational processes are initiated with the selection of a desired hand-movement trajectory and velocity profile. It is now common knowledge that a hand trajectory with a definite velocity profile does not define a unique pattern of joint rotations. An example of this effect is: when one reaches for an object with the hand and moves one's trunk forward at the same time, the hand trajectory remains invariant (Adamovich et al. 2001). But the arm's joint rotations are quite different. So the same trajectory is caused by several different patterns of component movements.

Hence, the computation of inverse dynamics of joint torques cannot take place unless the joint redundancy problem as described above is solved (for review, see Balasubramaniam & Feldman 2004). Moreover, a net joint torque does not define a unique force for each muscle crossing the joint, meaning that the inverse computation runs into an additional redundancy problem. Thus, from the point of initiation of the inverse computation a further redundancy problem is introduced and continues at each iterative level. Consequently, the nervous system faces an infinite regress of nested redundancy problems (Turvey 1990).

Multi-muscle redundancy: Just how much output can be programmed? This problem may be extended to redundancy at the level of the musculature as well. In just the same way that the trajectory does not map uniquely to the movement of the joints, muscle force does not determine a unique pattern of motor-unit recruitment. Inverse dynamical computational strategies exist with regard to the redundancy problem arising in the computations of individual muscle torques (Zajac et al. 2002). Although a variety of optimization criteria were used in the Zajac et al. study, it was concluded that because of the pattern of torques produced by multi-articular muscles, the inverse computations may fail to find the contributions of individual muscles.

For a complete and thorough model, it would be necessary to resolve the manner in which input signals to individual motor neurons (post-synaptic potentials) are computed to produce the desirable EMG output. Further, fundamental nonlinearities in the properties of motor neurons (such as threshold and plateau potentials) cannot be reversed without substantial simplifications of the dynamical input/output relationships in the system, which would reduce the reliability of model-based computations (Ostry & Feldman 2003).

Alternatives to emulation-based theories. Interesting alternatives to emulation-based approaches exist in which the problem of redundancy is treated fairly. For example, equilibrium-point approaches (Feldman & Levin 1995), uncontrolled manifold approaches (Scholz et al. 2000), and dynamical systems approaches (Turvey 1990). The fundamental difference between these approaches and emulation models is that motor output or behavior in the former is treated as an emergent property. In particular, according to Balasubramaniam and Feldman (2004), control neural levels may guide movement without redundancy problems only by predetermining in a task-specific way where, in spatial coordinates, neuromuscular elements may work, without instructing them how they should work to reach the desired motor output. Thus, no specific computations of the output are required – it emerges from interactions of the neuromuscular elements between themselves and the environment within the limits determined by external and control constraints.

Issues of implementation matter for representation

Francisco Calvo Garzón

Department of Philosophy, University of Murcia, Facultad de Filosofía, Edif. Luis Vives, Campus de Espinardo, Murcia, 30100, Spain. **fjcalvo@um.es**

Abstract: I argue that a dynamical framing of the emulation theory of representation may be at odds with its articulation in Grush's information-processing terms. An architectural constraint implicit in the emulation theory may have consequences not envisaged in the target article. In my view, "how the emulator manages to implement the forward mapping" is pivotal with regard to whether we have an emulation theory of representation, as opposed to an emulation theory of (mere) applied forces.

A dynamical framing of the emulation theory of representation, I contend, may be at odds with its articulation in Grush's information-processing terms. In my view, "*how* the emulator manages to implement the forward mapping" (sect. 2.2, para. 3, emphasis in original) is pivotal with regard to whether we have an emulation theory of representation, as opposed to an emulation theory of (mere) applied forces. Current work on the dynamics of representation – the *dynamic field approach* (Spencer & Schöner 2003; see also Erlhagen & Schöner 2002 and the references therein) – may furnish the means to implement Grush's emulation theory. According to the dynamic field approach, information gets represented by exploiting the neuroscientific concept of activation in the metric space of a dynamic field. Enduring behavior in an environment subject to perturbations, for example, gets explained in terms of how "activation in the field goes from a stable resting state through an instability (bifurcation) into a new attractor state – the self-sustaining state" (Spencer & Schöner 2003, p. 404). In an activation field, stabilities (e.g., attractor states) and instabilities (e.g., bifurcations) can be generated by dynamically "monitoring and updating movements using sensory feedback" (Spencer & Schöner, p. 394).

The dynamic field approach and its use of activation states fit nicely with potential extensions of Grush's model. Damasio's (1994) theory of reason and emotion, for instance, could be cashed out in terms of (cognitive) dynamic simulations that make use of inhibitory competition. A Hopfield-like competitive dynamical network would account for the instabilities and states of attraction

https://doi.org/10.1017/S0140525X0422009X Published online by Cambridge University Press