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 that shape the evolution of the activation field. Granting this framework for argument's sake, however, an architectural constraint implicit in the emulation theory may have consequences not envisaged in the target article.

Grush favors an *articulated* reading of emulation such that behavior gets explained in terms of the dynamical interactions of the relevant state variables. It is noteworthy that the vast majority of cognitive scientists would agree that the content of these variables must allow for *discontinuities* (although see below). Mental activity differs from motor responses (Thelen et al. 2001) in that, unlike the case of the motor system, where states *always* change continuously, mental content need not evolve that way. An activation field for a mental task may show a decay of activity, say, at point A, and a subsequent peak at a different location B, without a continuous shift of activation at intermediate positions. Higher-level cognition allows for responses whose informational content does not relate in a systematic way to the informational content of otherwise similar responses. Put bluntly, not all systematic patterns of behavior can exploit exclusively the continuities in state space evolution, as is the case in the motor system approach. Bearing in mind Grush's ultimate goal of "addressing other psychological capacities such as reasoning, theory of mind, and language" (sect. 1, para. 3) within the emulation theory framework, the emulation theory of representation now faces a dilemma.

On the one hand, someone may wish to call into question the demand for discontinuities; elsewhere (Calvo Garzón, in preparation) I argue that we may not be able to spell out a general theory of cognition in dynamical terms while allowing for discontinuities. On the other hand, the emulation theory may try to exploit mathematical resources of the dynamic field approach that would permit the emulator to exploit discontinuities (see Spencer & Schöner 2003). In either case, we are in trouble. If the need for discontinuities is ignored, Grush may be obliged to favor a non-articulated reading of his theory; a reading that should still allow us to account in computational terms for complex features such as recursion. Unfortunately, the "lookup table" option does not seem very attractive. For one thing, neurobiological evidence (O'Reilly & Munakata 2000) tells us that memory is not likely to deliver the goods, ecologically speaking, by implementing lookup tables.

Exploring non-articulated options, nevertheless, would take us far afield and, since Grush himself favors the articulated reading, we may for present purposes agree with him and ignore non-articulated alternatives. In any case, one might argue, the emulator theory of representation may be easily reconciled with the employment of discontinuities. According to Grush, what "allows us to acknowledge the action/behavioral bias of perception *without* becoming anti-representationalists about perception" (sect. 5.4, para. 4, emphasis added) is the coupling of cognitive agents with their surrounding environment. His model emulates the interactions that take place in contexts of *situated* cognition. Someone may wonder whether such acknowledgment is straightforwardly compatible with the positing of discontinuities. But we need not press further in that direction. It is regrettable that the discontinuous (dynamic field) use of the term "representation," however it gets fleshed out ultimately, is metaphysically weightless. It refers to the uncontroversial fact that sensory inputs get transformed into neural output. Such an approach, I contend, is compatible with an applied-force interpretation of emulation theory.

It is my hypothesis that a (dynamic systems theory) continuous and situated approach can synthesize different models of higher-level as well as lower-level cognition at the expense of having to eschew, rather than revise, the (computationalist) function-approximator approach that is explicitly endorsed in the target article. We may need to zoom back to enlarge the picture, and turn to questions concerning the role played by the information-processing paradigm and the role that potential contenders may play in the future. It is fair to say, nonetheless, that the fact that Grush's theory falls neatly within the information-processing paradigm does not mean that the above problems are insurmountable. Grush may be able to explain the evolution of the states of the system in terms of the predictions generated for all possible state variables while remaining representationalist. But he needs to say how. Issues of implementation do matter.

ACKNOWLEDGMENT
The author was supported by a Ramón y Cajal research contract (Spanish Ministry of Science and Technology) during the preparation of this commentary. I thank Esther Thelen for helpful discussions.

Testable corollaries, a conceptual error, and neural correlates of Grush's synthesis

Thomas G. Campbell and John D. Pettigrew
The Vision Touch Hearing Research Centre, The University of Queensland, Brisbane, 4072 Australia. **tomcampbell@uq.edu.au**
j.pettigrew@mailbox.uq.edu.au
http://www.vthrc.uq.edu.au/tomc/tomc.html
http://www.uq.edu.au/nuq/jack/jack.html

Abstract: As fundamental researchers in the neuroethology of efference copy, we were stimulated by Grush's bold and original synthesis. In the following critique, we draw attention to ways in which it might be tested in the future, we point out an avoidable conceptual error concerning emulation that Grush seems to share with other workers in the field, and we raise questions about the neural correlates of Grush's schemata that might be probed by neurophysiologists.

1. Testable corollaries. Grush presents a new synthesis that unites motor control, visual imagery, and perception under a single rubric. This bold, integrative step has a number of testable corollaries. For example, if these three seemingly distinct systems share the same underlying neural mechanisms, then it follows that they must also share a common timing mechanism. This point was presciently put forward by the physicist Richard Feynman, who talked of the need for the brain to have a "master clock" (Feynman 2001). One would therefore expect to find a common timing mechanism that links visual perception and motor control. Preliminary evidence for such a surprising link has recently been provided (Campbell et al. 2003).

Another specific example where predictions of Grush's schema can be explicitly tested is in the "mirror neuron" system, whose beautiful exposition in premotor cortex we owe to Rizzolatti and colleagues (Rizzolatti et al. 1999). One of the major puzzles that is presented by this work – whose lack of suggested correlations with the major components of Grush's emulators is perhaps its greatest weakness – could both be illuminated by Grush's approach and in turn help make explicit predictions on neural systems. The puzzle is the following: How does a neural system that has been set up to encode a specific, complex motor act also know how that act's performance will appear to an outside observer?

The extreme specificity shown by mirror neurons makes it highly unlikely that this outcome is the result of coincidental experience (the view that that brain is plastic porridge and that all can be explained by experience-dependent plasticity). Instead, it seems more likely that visual perception and motor performance share a common organizational structure, as Grush proposes, that is responsible for the surprising correspondence between the motor and visual (and even auditory) properties of the mirror neuron. We find it difficult to escape the conclusion from these considerations that even basic aspects of visual perception must have a strong "efferent" aspect. This could be tested explicitly using the predictions of Grush's formulation in the context of the mirror-neuron system.

2. A conceptual error about efference copy. In his synthesis Grush uses the term "efference copy" as a synonym of corollary discharge. We believe that this blurring of the distinction represents an unhelpful oversimplification of the corollary discharge efference copy (CDEC) system. Although this point may seem to be only semantic, we believe that a recursive error is generated by the