

The emulation theory of representation: Motor control, imagery, and perception

Rick Grush

Department of Philosophy, University of California, San Diego,
La Jolla, CA 92093-0119

rick@mind.ucsd.edu <http://mind.ucsd.edu>

Abstract: The *emulation theory of representation* is developed and explored as a framework that can revealingly synthesize a wide variety of representational functions of the brain. The framework is based on constructs from control theory (forward models) and signal processing (Kalman filters). The idea is that in addition to simply engaging with the body and environment, the brain constructs neural circuits that act as models of the body and environment. During overt sensorimotor engagement, these models are driven by efference copies in parallel with the body and environment, in order to provide expectations of the sensory feedback, and to enhance and process sensory information. These models can also be run off-line in order to produce imagery, estimate outcomes of different actions, and evaluate and develop motor plans. The framework is initially developed within the context of motor control, where it has been shown that inner models running in parallel with the body can reduce the effects of feedback delay problems. The same mechanisms can account for motor imagery as the off-line driving of the emulator via efference copies. The framework is extended to account for visual imagery as the off-line driving of an emulator of the motor-visual loop. I also show how such systems can provide for amodal spatial imagery. Perception, including visual perception, results from such models being used to form expectations of, and to interpret, sensory input. I close by briefly outlining other cognitive functions that might also be synthesized within this framework, including reasoning, theory of mind phenomena, and language.

Keywords: efference copies; emulation theory of representation; forward models; Kalman filters; motor control; motor imagery; perception; visual imagery

1. Introduction

The idea that one of the central tasks performed by the brain is to internally model various brain-external elements and processes is not new. In the twentieth century, Kenneth Craik was one of the most explicit proponents of this view (Craik 1943). In various guises the view has been taken up as an approach to understanding reasoning (Johnson-Laird 1983), theory of mind phenomena (Gordon 1986), mental imagery (Kosslyn 1994), and even aspects of motor control (Kawato 1999). The metaphor of “internal modeling” aside, these approaches have not (yet) been synthesized into anything approaching a unified and flexible framework.

In this article I have four goals. The first is to articulate an information-processing strategy that I will call the *emulation theory of representation*. This strategy is developed using tools from control theory and signal processing, especially drawing on pseudo-closed-loop control and Kalman filtering. I will try to use just enough mathematical formalism to ensure that the main ideas are clear. The second goal is to show the use of this framework in understanding certain aspects of motor control and motor imagery. Many researchers in the two fields of motor control and motor imagery currently appeal to constructs related to those I will develop, but such appeals rarely go into much detail concerning the overall information-processing structures involved, and little by way of clear synthesis has

emerged. In providing such a structure, I hope to do a service to both of these areas of research by providing a framework within which various results can be synthesized, and within which a number of issues can be more clearly stated so as to avoid certain kinds of errors.

The emulation framework I will articulate not only allows for a great deal of motor control and motor imagery work to be synthesized in interesting ways, but it also synthesizes certain aspects of visual imagery and visual perception as well. Outlining such a synthesis is the third goal. The final goal, addressed in the last section, is to briefly explore the prospects for addressing other psychological capacities such as reasoning, theory of mind, and language, within the same framework.

RICK GRUSH is Associate Professor of Philosophy at the University of California, San Diego, the same university at which he received a joint doctorate in cognitive science and philosophy in 1995. His research in theoretical cognitive neuroscience is directed at understanding the nature of representation and the structure and neural implementation of fundamental representational capacities of nervous systems, including space, time, and objects. These topics will be explored in detail in his in-progress book, *The Machinery of Mindedness*.




Figure 1. A simple feed-forward control schematic. “Plant” is the control theory term for the controlled system; in the case of motor control, the plant is the body, specifically the musculoskeletal system (MSS) and relevant proprioceptive/kinesthetic systems. Here, the “plant” includes the MSS, sensors, and any random perturbations, such as muscle twitches and noise in the sensors. These components will be treated separately in later sections.

2. Motor control: Forward models and Kalman filters

2.1. Feed-forward and feedback control

The nature of the interaction between the motor centers of the brain and feedback from the body during fast, goal-directed movements has been a long-standing controversy in motor control (Desmurget & Grafton 2000; van der Meulen et al. 1990). On one side are those who claim that the movements are ballistic or *feed-forward*, meaning roughly that the motor centers determine and produce the entire motor sequence (sequence of neural impulses to be sent down the spinal cord) on the basis of information about the current and goal body configurations; see Figure 1. As a result of this motor volley, the body then moves to a configuration near the goal state. It is only at the very end of the movement – when fine adjustments are required – that visual and proprioceptive/kinesthetic feedback are used; the bulk of the motor sequence is determined and executed without feedback.

On the other side of the debate have been those who argue for *feedback* control. In the form most opposed to feed-forward control it is claimed that though there is a goal, there is no motor *plan* prior to movement onset. Rather, the motor centers continually compare the goal configuration to the current configuration (information about which is provided through vision or proprioception/kinesthesia), and simply move the current configuration so as to reduce the difference between it and the goal configuration. A simplified schematic for this sort of feedback control is shown in Figure 2.

In both schemes, the control process breaks down into two components, the *inverse mapping* and the *forward mapping*. The “forward” in forward mapping is meant to capture the fact that this part of the process operates in the

direction of causal influence. It is a mapping from current states and motor commands to the future states that will result when those motor commands exert their influence. Clearly, this is the mapping implemented by the actual musculoskeletal system, which moves on the basis of motor commands from its current state to future states. On the other hand, the controller implements the *inverse* mapping. It takes as input a specification of the future (goal) state, and determines as output what motor commands will be required in order to achieve that state. This mapping is just the inverse (or more typically, *an* inverse) of the forward mapping. Hence, when placed in series, a good controller and plant form an *identity mapping*, from goal states to goal states. See Figure 3.

Where the schemes differ is on *how* the controller implements the inverse mapping to produce the control sequence. In the feed-forward scheme, the entire motor plan is determined largely before movement onset. This often requires a good deal of sophistication on the part of the controller. In the feedback scheme, the control sequence emerges over time as the process of interaction between the controller and plant unfolds. But in both cases, the controller produces a motor control sequence, and the body executes it to move into the goal configuration. (For more on these two schemes, especially the kinds of evidence historically used to argue in favor of each, see Desmurget & Grafton 2000.)

2.2. Emulators (forward models)

But feed-forward and feedback control do not exhaust the alternatives. There has been growing recognition among researchers in motor control that schemes involving the use of *forward models* of the musculoskeletal system (henceforth MSS) are promising (Kawato 1999; Wolpert et al. 2001). I will use “emulator” as a more descriptive synonym

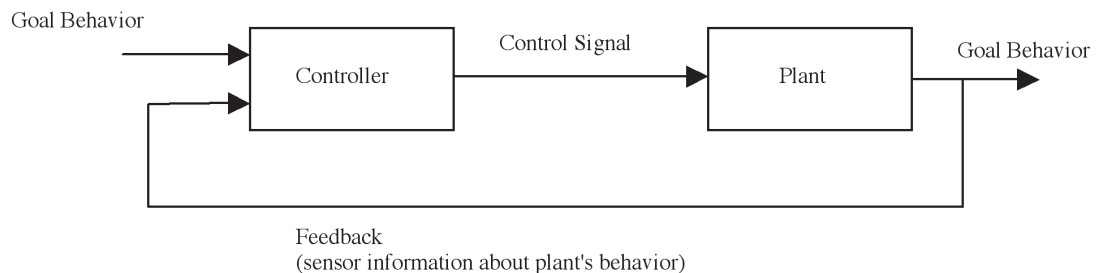


Figure 2. A simple feedback control schematic. Sensors (a component of the plant) measure critical parameters of the plant, and this information is continually provided to the controller. The controller uses this feedback in its evolving production of the control signal.

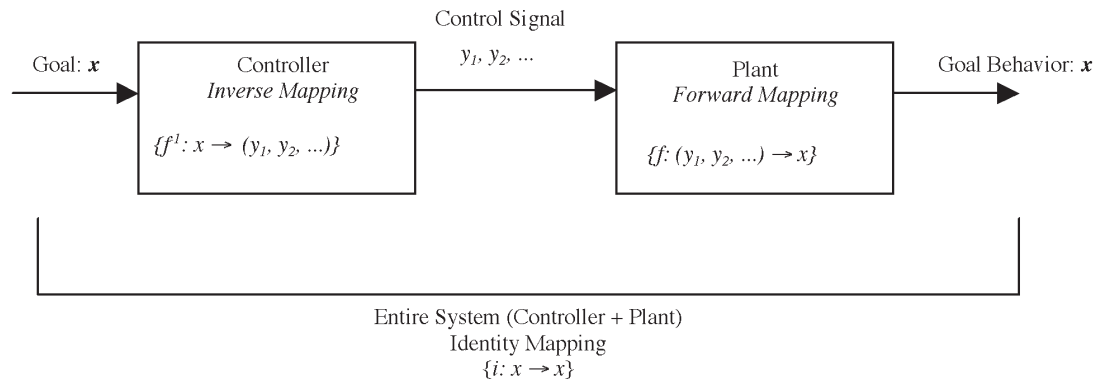


Figure 3. Forward and inverse mappings.

for “forward model.” In such schemes, the controller exploits a continuing stream of feedback provided by an emulator of MSS dynamics, which is driven by efference copies of the motor commands sent to the body. The simplest such scheme is shown in Figure 4.

In this scheme the controlling system includes not just the controller, but also an emulator. The emulator is simply a device that implements the same (or very close) input-output function as the plant. So when it receives a copy of the control signal (it is thus getting the same input as the plant), it produces an output signal, the emulator feedback, identical or similar to the feedback signal produced by the plant. Such feedback has a number of uses. One use is that emulator feedback can be subject to less delay than feedback from the periphery (Desmurget & Grafton 2000; Wolpert et al. 1995). This will be explored in more detail in section 2.5.

There are two points to highlight that will be important later on. First, for these specific purposes it does not matter *how* the emulator manages to implement the forward mapping (though of course for other purposes it matters quite a bit, and empirical evidence is relevant to deciding between the options I will discuss below). One possibility is

that an emulator might simply be a large associative memory implementing a lookup table whose entries are previously observed musculoskeletal input-output sequences; and upon receiving a new input, it finds the closest associated output (see Miles & Rogers 1993). Another possibility is for the emulator to be what I will call an *articulated model*. The *real* MSS behaves the way it does because it has a number of state variables (such as elbow angle, arm angular inertia, tension on quadriceps) that interact according to the laws of dynamics and mechanics. Some of these variables are measured by, for example, stretch receptors and Golgi tendon organs. This measurement constitutes bodily proprioception and kinesthesia: the “feedback” in control theoretic terms. Similarly, an articulated emulator is a functional organization of components (*articulants*) such that for each significant variable of the MSS, there is a corresponding articulant, and these articulants’ interaction is analogous to the interaction between the variables of the MSS. For example, there would be a group of neurons whose firing frequency corresponds to elbow angle; and this group makes excitatory connections on another group that corresponds to arm angular inertia, such that, just as an

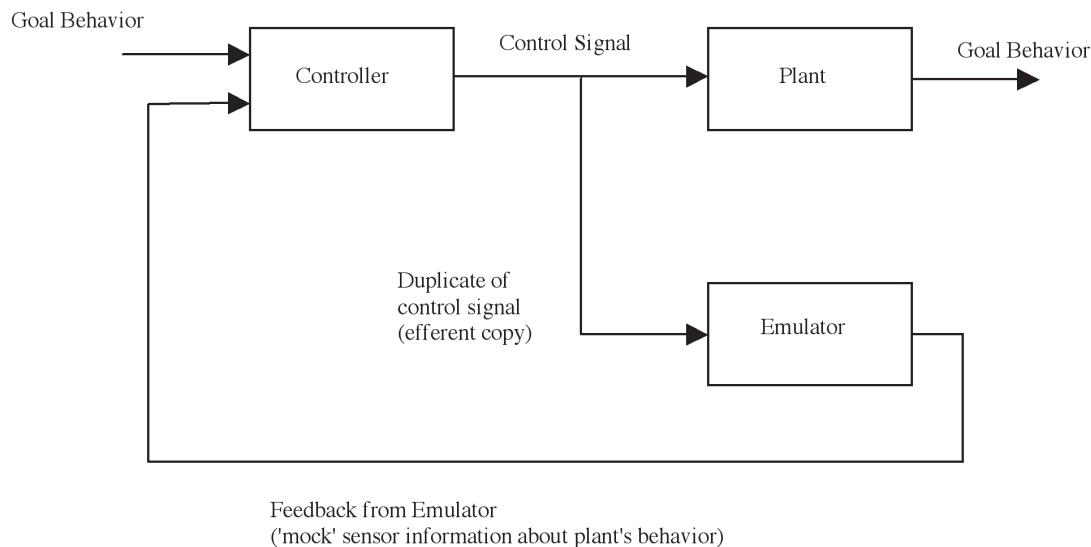


Figure 4. A simple pseudo-closed-loop control schematic (Ito 1970; 1984). A copy of the control signal – an *efference copy* – is sent to a subsystem – an emulator – that mimics the input-output operation of the plant. Because the emulator is given the same input as the plant (its input is a copy of the signal sent to the plant) it produces a similar output.

increase in elbow angle results in an increase in arm angular inertia, an increase in the firing rate of the first group of neurons instigates an increase in the firing rate of the second (see Grush 1995; Kawato et al. 1987). And just as the real MSS is subject to a measurement that provides proprioceptive and kinesthetic information, the articulated emulator can have a “measurement” taken of the same variables, and thus yield a mock sensory signal. There are more than just these two options, but for purposes of this article something more along the lines of the articulated emulator will be assumed – though this is an empirical issue, of course.

The second point is that emulators, whether lookup tables or articulated models, must, under normal conditions, have a certain degree of plasticity. This is because the systems they emulate often alter their input-output function over time. This phenomenon is *plant drift* – in the case of mechanical plants, belts loosen, gears wear, some parts get replaced by others that are not exactly the same; in the case of the body, limbs grow, muscles get stronger or weaker over time. Whatever the details, a particular motor command might lead to one MSS output at one time, but lead to a slightly different output at some time months or years later. In order to remain useful, the overall control system needs to monitor the input-output operation of the plant, and be able to slowly adjust the emulator’s operation so as to follow the plant’s input-output function as it drifts.

2.3. Kalman filters

The advantage of pseudo-closed-loop control is that it is conceptually simple, making it an easy way to introduce certain ideas. It is too simple as it stands to be of much use in explaining real biological systems. The next conceptual in-

crement needed in order to bring these basic ideas up to speed is the Kalman filter (which I will abbreviate KF). My discussion of KFs here will make a number of omissions and simplifications. For example, I discuss only state estimation/updating and not variance estimation/updating; I discuss only discrete linear models and ignore generalizations to continuous and nonlinear systems. My goal is simply to introduce those aspects of KFs that are important for the remaining discussion.¹ The technique, a standard version of it anyway, is diagrammed in Figure 5.

First, we need a description of the problem to be solved by the KF. The problem is represented by the top part of the diagram, within the dashed-line box. We start with a *process* consisting of a system of k state variables, and whose state at time t can thus be described by the $k \times 1$ vector $r(t)$. The process’ state evolves over time under three influences: first, the process’ own dynamic (represented here by the matrix V); second, process noise, which is any unpredictable external influence; and third, the driving force, which is any predictable external influence. Without the noise or the driving force, the process’ state at any given time would be a function of its state at the previous time: $r(t) = Vr(t-1)$, where V is a $k \times k$ matrix that maps values of r into new values of r . Noisiness (random perturbations) in the evolution of the process can be represented by a small, zero-mean, time-dependent $k \times 1$ vector $n(t)$, and the driving force as another $k \times 1$ vector $e(t)$. Thus, the process at time t is: $r(t) = Vr(t-1) + n(t) + e(t)$.

The real signal $I(t)$ is a *measurement* of states of this process. We can represent this measurement as a $k \times h$ measurement matrix O , that maps $r(t)$ to the $h \times 1$ signal vector $I(t)$: $I(t) = Or(t)$. (An obvious special case is where O is the identity matrix I , in which case $I(t) = r(t)$.) We can

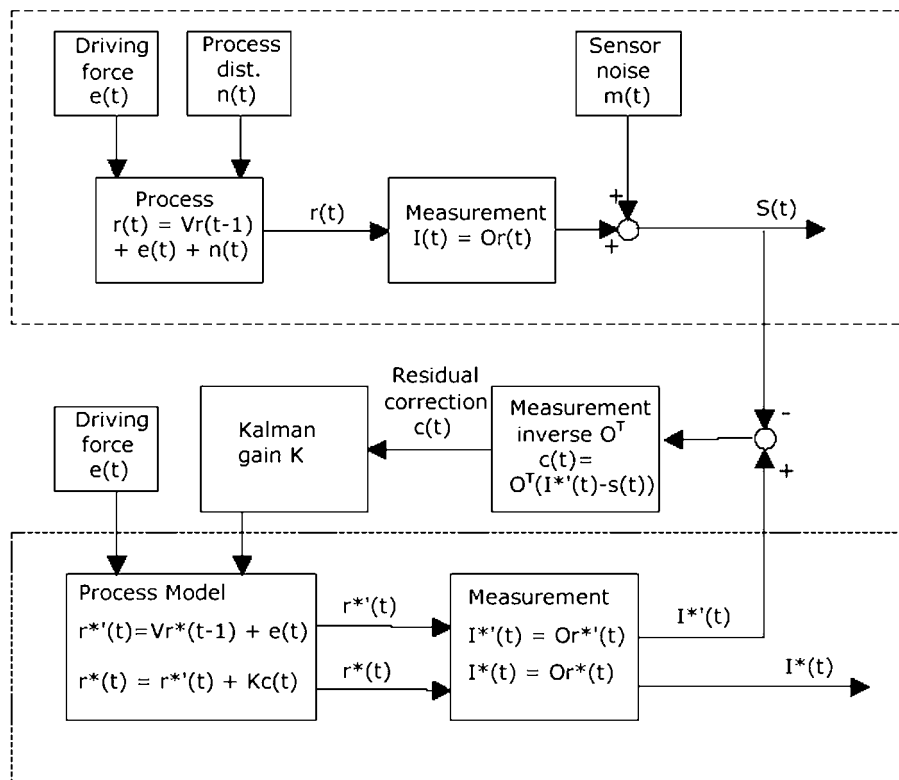


Figure 5. A basic Kalman filtering scheme.

represent unpredictable noise in the measurement process – the sensor noise – as another small time-dependent zero-mean $h \times I$ vector $m(t)$, and so the actual output, the *observed signal*, is $S(t) = I(t) + m(t)$. The problem to be solved is to filter the sensor noise $m(t)$ from the observed signal $S(t)$ in order to determine what the real, noise-free, signal $I(t)$ is.

The core idea behind the KF is that it maintains an optimal estimate of the real process' state, and then subjects this estimate to the same measurement that produces the real signal from the real process. The result is an optimal estimate of the noise-free signal. The KF's estimate of the process' state is embodied in the state of a *process model*, an articulated model of the process. We can represent the process model's state as $r^*(t)$. The KF keeps $r^*(t)$ as close as it can to $r(t)$, meaning that it tries to match, as closely as possible, the value of each of the k state variables of r^* to the corresponding state variable of r .

So far it should be clear how, with the benefit of an accurate state estimate $r^*(t)$, the KF can produce an optimal estimate of the real signal $I(t)$. Now we must examine how it maintains an optimal state estimate $r^*(t)$. This is done in two steps. The first step is often called the *time update*. Given the previous state estimate $r^*(t-1)$, the KF produces an expectation or prediction of what the state at t will be by evolving $r^*(t-1)$ according to the same dynamic V that governs the evolution of the real process and adding the driving force $e(t)$.² The result of this time update is the *a priori* estimate, $r^{*'}(t)$ (note the prime), and as stated, it is arrived at thus: $r^{*'}(t) = Vr^*(t-1) + e(t)$. It is called the *a priori* estimate because it is the estimate arrived at before taking information from the observed signal into account. Qualitatively, the KF says, "Given my best estimate for the previous state and given how these states change over time, and also given my knowledge of the driving force, what should I expect the next state to be?" This is the *a priori* estimate.

The next step, often called the *measurement update*, uses information from the observed signal to apply a correction to the *a priori* estimate. This is done in the following way (roughly – again I must note that my description here is making a number of short cuts): The *a priori* estimate $r^{*'}(t)$ is measured to produce an *a priori* signal estimate $I^{*'}(t)$. This is compared to the observed signal $S(t)$. The difference is called the *sensory residual*. From this residual it is determined how much the *a priori* estimate would have to be changed in order to eliminate the residual altogether. This is done by pushing the residual through the inverse of the measurement matrix O , O^T . This is the *residual correction* $c(t)$. Though the KF now knows exactly how to alter the *a priori* estimate in order to eliminate the sensory residual, it typically does not apply the entire residual correction. Why? The residual correction is how much the *a priori* state estimate would have to be altered to eliminate the sensory residual. But not all of the sensory residual is a result of the inaccuracy of the *a priori* estimate. Some of it is a result of sensor noise. Thus, the *a priori* estimates $r^{*'}(t)$ and $I^{*'}(t)$ might be very accurate, and the sensory residual due mostly to the sensor noise.

The KF determines how much of this correction should actually be applied based on the KF's estimates of the relative reliability of the *a priori* estimate versus the noisy observed signal $S(t)$. (The determination of the relative reliability is part of the process I have not gone into, it is the determination of the *Kalman gain*.) To the extent that the

process noise is small compared to the sensor noise, the *a priori* estimate will be more reliable than the observed signal, and so a smaller portion of the residual correction is applied to the *a priori* estimate. To the extent that the sensor noise is small compared to the process noise, the observed signal is more reliable than the *a priori* estimate, and so a greater portion of the residual correction is applied.

Qualitatively, the KF compares its expectation of what the signal should be to what it actually is, and on the basis of the mismatch adjusts its estimate of what state the real process is in. In some conditions, such as when the sensors are expected to be unreliable, the expectation is given more weight than the signal. In other conditions, such as when the process is less predictable but sensor information is good, the expectation is given less weight than the signal.

The result of the measurement update is the *a posteriori* estimate $r^*(t)$, which is a function both of the expectation and the observation. This estimate is measured using O in order to get the final estimate $I^*(t)$ of the noise-free signal $I(t)$.

2.4. Kalman filtering and control

While KFs are not essentially connected to control contexts, they can be easily incorporated into control systems. Figure 6 shows a Kalman filter incorporated into a control scheme very similar to the pseudo-closed-loop scheme of Figure 4. In Figure 6, everything within the dotted-line box is just the KF as described in the previous section, and shown in Figure 5. The only difference is that the external driving force just is the control signal. Everything in the dashed-line box is functionally equivalent to the plant in the pseudo-closed-loop scheme of Figure 4, and everything in the dotted-line box is functionally equivalent to the emulator. The box labeled "plant" in Figures 1–4 did not separate out the process, sensors (a.k.a. measurement), and noise, but lumped them all together. The "emulator" box in Figure 4 similarly lumped together the emulation of the process and the emulation of the measurement via sensors.

In effect, this is a control scheme in which an articulated emulator is used as part of a Kalman filter for the purpose of filtering noise from the plant's feedback signal.

Note that the scheme in Figure 6 subsumes closed-loop control and pseudo-closed-loop control as degenerate cases. If the Kalman gain is set so that the entire sensory residual is always applied, the scheme becomes functionally equivalent to closed-loop control. When the entire sensory residual is applied, the *a posteriori* estimate becomes whatever it takes to ensure that $I^*(t)$ exactly matches $S(t)$. Thus, the signal sent back to the controller will always be whatever signal actually is observed from the process/plant, just as in closed-loop control.

On the other hand, if the Kalman gain is set so that none of the residual correction is applied, then the *a priori* estimate is never adjusted on the basis of feedback from the process/plant. The state estimate of the process model (emulator) evolves over time exclusively under the influence of its own inner dynamic and the controller's efference copies, and the feedback sent to the controller is just a current measurement of this encapsulated estimate, just as in pseudo-closed-loop control.

I will refer to systems like that in Figure 6 as KF-control schemes. At the same time, I will indulge in a certain degree of flexibility in that I will take it that extensions to con-

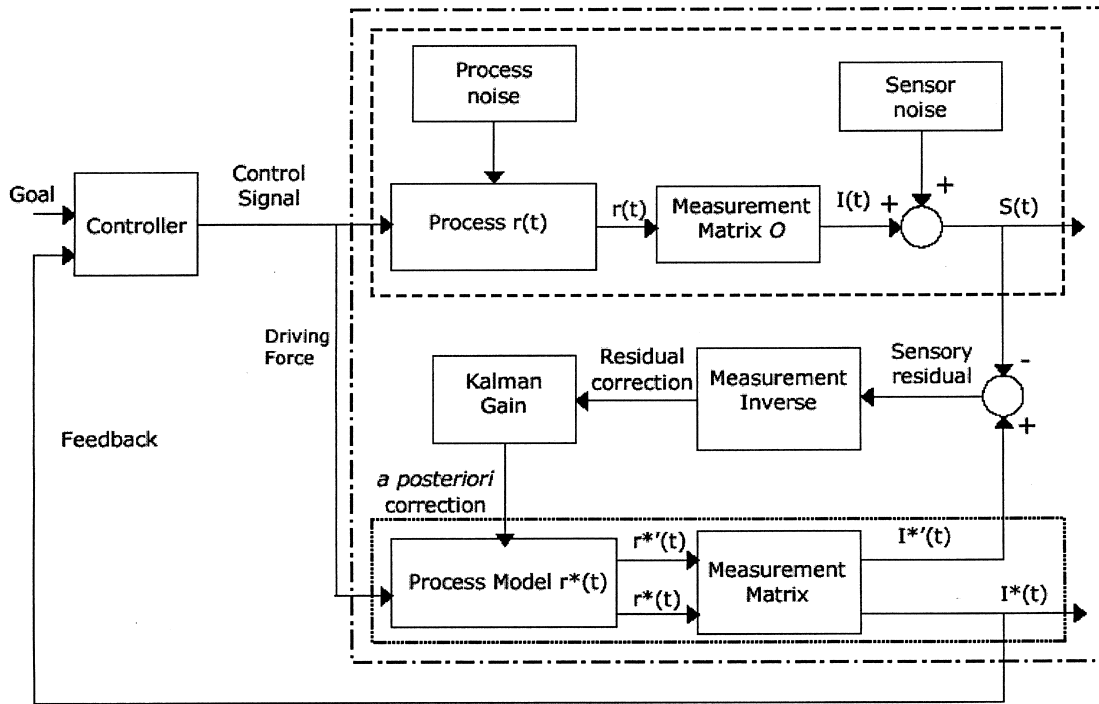


Figure 6. A control scheme blending pseudo-closed-loop control and a Kalman filter.

tinuous and nonlinear systems are included (see the references in Note 1 for more on these cases), and that the operation may not always be optimal, as when the dictates of the Kalman gain are overridden in order to produce imagery (see sect. 3).

Before closing this section, it might help to provide a qualitative example of a problem structure and solution structure that take the form described by the formalism above. The example is ship navigation.³ The “process” is the ship and its location, which at any time is a function of its location at the previous time, plus the driving force (the engine speed, rudder angle), plus the “process noise” (unpredictable winds and ocean currents). The “signal” in this case is bearing measurements, which consist *inter alia* of the observed angles to known landmarks or stars. The observation-measurement process is not perfect, and this imperfection can be effectively represented as “sensor noise” added to an ideal bearing measurement. The task – exactly the same as the Kalman filter’s task as described in section 2.3 – is to determine, on the basis of the imprecise bearing measurements, the ship’s actual position (or equivalently, what the noise-free, or perfect, bearing measurements would be). To solve this, the navigation team maintains a model of the process, implemented on a map. At each time step (or “fix cycle”), the team projects where the ship should be, on the basis of the estimate from the previous fix cycle and the known driving force. This a priori prediction is then combined with the actual observations in order to determine a “best guess” region within which the ship is located: the a posteriori estimate. And this a posteriori estimate becomes the initial estimate for the next cycle. And, just as in the KF, if an ideal navigator knows that there are no unpredictable winds or currents, but the bearing measurements are very imprecise, then the prediction based on the previous location – the a priori estimate – will be weighted more heavily than the bearing measurements (low Kalman

gain). If, on the other hand, the bearing measurements are quite reliable but there are a lot of winds and currents (high process noise), then the “sensory residual” will be weighted more heavily (high Kalman gain). Of course, real navigators do not calculate anything like a Kalman gain matrix for optimally combining a priori expectation and observed signals. Real navigators are often good, but seldom optimal. This is one of the many ways in which the emulation framework relaxes the strict requirements of the Kalman filter, while maintaining the general information-processing structure.

The general structure of the problem and the general structure of the KF’s solution – a structure captured by the emulation framework – are quite unmysterious. One system (a ship’s crew, a brain) is interacting with another system (a ship, an environment, a body) such that the general principles of how this system functions are known, but the system’s state is not *entirely* predictable owing to process noise (unpredictable currents, bodily or environmental perturbations), and imperfect sensors. The solution is to maintain a model of the process – done by a part of the ship’s crew, the navigation team; or specialized emulator circuits in the brain – in order to provide predictions about what its state will be; and to use this prediction in combination with sensor information in order to maintain a good estimate of the actual state of the system that is being interacted with – an estimate that makes use of sensor information, but is not limited by the many sub-optimality of the bare sensory process.

2.5. Evidence for the emulation framework

The emulation framework is a specification of a certain kind of information-processing structure. It is my hypothesis that the brain uses this information-processing structure for a number of tasks, and it is also my contention that this hypothesis, if true, can serve to synthesize a wide range of oth-

erwise seemingly unrelated models and results. My evidence for the claim that the brain uses the emulation framework is distributed throughout this paper, but in most cases it takes the following form. First, I introduce some specific area of potential application, such as motor control or visual imagery. I then pick out some exemplar model or result, already existing in the cognitive neuroscientific literature in this particular domain, which offers a particularly good example of how the emulation framework applies in that domain. Because the exemplar model is a specific application of the broader synthesizing emulation framework, the evidence, typically produced by the researchers who proposed the specific model, becomes evidence for the emulation framework's application to that domain. And vice versa, any evidence that would count against the specific exemplar model in question would, *prima facie*, count as evidence against the emulation framework's applicability to that domain. Since it is the goal of this article to articulate the synthesizing framework, which requires articulating the framework and applying it to a number of domains, space limitations prevent me from going into anything approaching a convincing level of detail concerning the evidence for the model in each of the areas covered. Rather, my goal is to cover just enough of the evidence, typically in the form of one or two exemplar models or experimental results, to make the character of the proposed synthesis clear – to give specific flesh to not only how the emulation framework applies to that domain, but also to indicate the kinds of evidence that count for and against the emulation framework in that domain. For those who wish to see more detailed evidence within each domain of application, I will provide references to specific studies and models, done by specialists within the specific field.

2.6. Motor control

The first application is within the domain of motor control, and the exemplar of an application of the emulation framework to motor control I use here is a model by Wolpert, Ghahramani, and Jordan (Wolpert et al. 1995). The goal of their project was to provide a model that would explain the information-processing mechanisms used by the nervous system to maintain an estimate of the body's state during movements. Wolpert et al. distinguished three possibilities: pure peripheral sensory information, pure centrally generated predictions based on motor outflow (a.k.a. knowledge of the driving force via efference copies), and a combination of these two integrated via a Kalman filter scheme essentially identical to that described in section 2.4. The data were subjects' estimates of the position of their own hands after movements of varying lengths of time executed without visual feedback. The observed pattern demonstrated that as the duration of the movement increased from .5 second (the shortest movement duration) to 1 second, the magnitude of subjects' over-estimation of the position of their hands increased. But after a maximum at about one second, the magnitude of the overestimation decreased. This same temporal pattern was observed in three conditions: assistive force, resistive force, and no external force.

Wolpert et al. argued that neither the pure sensory inflow nor the pure motor outflow-based models could account for this pattern, but that the Kalman filter model could. Because proprioceptive and kinesthetic feedback is not available during the initial stages of the movement, the state es-

timate early on is based almost entirely on the uncorrected predictions of the process model. Initial error (for which there is an initial overestimate bias) in this estimate compounds as time progresses in absence of correction from sensory signals. However, as feedback becomes available from the periphery, this information is used to make corrections to the state estimate that begin to show their influence about one second after movement onset. Hence, there is a drop in the magnitude of the overestimation after about one second.

Many models and experimental results published in recent years provide strong evidence in support of the idea that at least some aspects of motor control are subserved by some kind of KF information-processing structure. These include Blakemore et al. (1998), Kawato (1999), Wolpert et al. (2001), Krakauer et al. (1999), and Houk et al. (1990). Mehta and Schaal (2002) study the dynamic characteristics of subjects' movements in the pole balancing task (including "blackout" trials where visual feedback is blocked for some brief period), and show that the data strongly suggest that a Kalman filter-like mechanism is used by the central nervous system to process sensory information during the execution of this task. As for what evidence would count against the emulation framework's applicability to motor control, it is whatever evidence would count against the models just cited. For example, a lack of a time-varying estimation pattern as found by Wolpert et al. (1995), or behavioral patterns of pole balancing tasks as described in Mehta and Schaal (2002), which would disconfirm a Kalman filter mechanism and implicate instead any of the other possibilities they claim are disconfirmed.

3. Motor imagery

3.1. The emulation theory of motor imagery

The Kalman gain determines the extent to which the sensory residual influences the a priori estimate produced by the emulator – qualitatively, the degree to which raw sensory input trumps or does not trump expectation. Typically, the Kalman gain is determined on the basis of the relative *variance* of the prediction and the sensor signal errors. This is part of what allows the KF to be an optimal state estimator, and in control contexts having optimal information is often good. The Kalman gain allows us to breathe some much-needed flexibility and content into the stale and overly metaphorical distinction between top-down and bottom-up sensory/perceptual processing. In the terminology developed in the previous section, we can see that a KF processor is top-down to the extent that the Kalman gain is low – the lower the Kalman gain, the more the final percept is determined by the expectation, which in turn is a function of the operation of the model's inner dynamic as driven by efference copies, if any. The higher the Kalman gain, the more this inner dynamic is overridden by sensory deliverances. That is, the same system not only implements both sorts of processing, but can flexibly and optimally combine them as conditions and context dictate.

Section 5 will explore this in the context of perception. For now, I want to draw attention to the fact that a system that is set up to allow for flexibility in this regard can be exploited for other purposes. Specifically, it can be used to produce *imagery*. Two things are required. First, the Kalman gain must be set so that real sensory information

has *no* effect. The emulator's state is allowed to evolve according to its own dynamic and the efference copies, if any; there is no "correction" or alteration from the senses. Second, the motor command must be suppressed from operating on the body.⁴ Therefore, on this view, motor imagery results from the control centers of the brain driving an emulator of the body, with the normal efferent flow disengaged from the musculature, and with the normal sensory inflow having no effect on the emulator's state and feedback. The very same process that is used to enhance overt motor control, as described in section 2.6, can thus be used to produce motor imagery. In this section I will briefly defend this emulation theory of motor imagery. The purpose is not only to give another example of the framework in operation, but also to show how it can help to clarify some of the issues within the field.

It will be helpful to compare the emulation theory of motor imagery to another proposal that superficially seems identical, but which is actually quite distinct. I have in mind the "simulation theory" of motor imagery, currently favored by a number of researchers (Jeannerod 2001; Johnson 2000a), according to which motor imagery is just the inner off-line operation of the efference motor centers of the brain. As Jeannerod and Frak put it, "motor imagery corresponds to a subliminal activation of the motor system" (Jeannerod & Frak 1999).

From the point of view of the emulation theory described above, the simulation theory is *half* correct. The part that is correct is that those areas corresponding to the controller – efferent motor areas – should be active during motor imagery. Accordingly, the evidence brought forward in favor of the *simulation* theory is evidence for at least half of the emulation theory. *The difference is that the simulation theory does not posit anything corresponding to an emulator*; as far as I can tell, the simulation theory is conceived against the backdrop of closed-loop control, and motor imagery hypothesized to be the free-spinning of the controller (motor centers) when disengaged from the plant (body). In the emulation theory, by contrast, imagery is not produced by the mere free-spinning operation of the efferent motor areas, but by the efferent motor areas driving an emulator of the musculoskeletal system. The next section quickly recaps two areas of evidence typically cited in support of the simulation theory. In section 3.3, I will discuss considerations favoring the emulation theory over the simulation theory. Before I move on, I should clarify something. I will use the terms "simulation theory" and "emulation theory" exactly as I have defined them in this paragraph. Given these definitions, it should be clear that the simulation theory and the emulation theory are not at all the same thing. They agree that the efferent motor centers are active during imagery. The simulation theory takes this by itself to be sufficient for motor imagery; the emulation theory does not, and claims that in addition, an emulator of the musculoskeletal system is needed and imagery is produced when the efferent motor centers drive this emulator. This distinction should be entirely obvious. To make an analogy: The emulation theory claims that motor imagery is like a pilot sitting in a flight simulator, and the pilot's efferent commands (hand and foot movements, etc.) are translated into faux "sensory" information (instrument readings, mock visual display) by the flight simulator which is essentially an emulator of an aircraft. The simulation theory claims that just a pilot, moving

her hands and feet around but driving neither a real aircraft nor a flight simulator, is sufficient for mock sensory information.

It may be the case that some of the researchers I cite as defending the simulation theory really mean something like the emulation theory, but just are not entirely clear about the distinction in their descriptions. If that is the case, then maybe even the defenders of the "simulation theory" are really agreeing with the emulation theory. In what follows I will compare the simulation theory as I describe it with the emulation theory as I describe it. If nobody really holds the simulation theory (as I define it), then at least some clarity will have been added to the literature on this topic.

3.2. The simulation theory of motor imagery

There are two related sorts of evidence cited by proponents of the simulation theory – the first is that many motor areas of the brain are active during motor imagery, the second concerns a number of isomorphisms between overt movements and their imagined counterparts.

That motor imagery involves the off-line activity of many motor areas is a widely replicated result (for a recent review, see Jeannerod 2001). PET studies of motor imagery consistently show selectively increased activity in premotor areas, supplementary motor areas, the cerebellum, and other areas. This is the defining feature of the simulation theory: that motor imagery is the psychological counterpart of the off-line operation of the efferent motor areas. Of the major motor areas canvassed in such studies, only primary motor cortex is (usually) conspicuously silent during motor imagery. This would seem to imply that the major signal bifurcation, where the efference copy is split from the "real" efferent signal, occurs just before primary motor cortex. (Though the implementation details are not entirely clear: Some studies, such as Richter et al. [2000], implicate primary motor cortex in motor imagery. Therefore, the functional separation of the efferent motor areas and the MSS emulator may not be easy to determine.)

Furthermore, a number of parallels between motor imagery and overt motor behavior have suggested to researchers that the two phenomena have overlapping neural and psychological bases (Deiber et al. 1998; Jeannerod & Frak 1999). To take a few examples, there are close relations between the time it takes subjects to actually perform an overt action and the time taken to imagine it (Jeannerod 1995). The frequency at which subjects are no longer able to overtly move their finger between two targets is the same for overt and imagined movement. There is evidence that even Fitts' Law is in effect in the domain of imagination (Decety & Jeannerod 1996).

Johnson (2000a) has provided compelling evidence that, when forming a motor plan, a subject's expectations about what range of grip orientations will be comfortable is very close to the range that actually is comfortable. Johnson argues that this result indicates not only that motor imagery is used for this task (it is motor rather than visual because the crucial factor is the biomechanics of the arm, not its visual presentation), but also that such imagery, precisely because it respects biomechanical constraints, is used to determine an effective motor plan before execution.⁵

The proponents of the simulation theory point out that

not only are motor areas active during motor imagery, but the isomorphisms between the observed activity of motor areas and the motor imagery suggest that motor imagery is in fact the product of the operation of motor centers, whose operational parameters are tuned to overt performance and hence recapitulated in covert performance.

3.3. Emulation versus simulation

None of the results appealed to in the previous section distinguish the simulation from the emulation theories, as both expect the controller (efferent motor centers) to be active during the production of motor imagery. The difference is that the emulation theory claims that mere operation of the motor centers is not enough; to produce imagery, the motor areas must be driving an emulator of the body (the MSS and relevant sensors).

There are reasons for preferring the emulation theory – for thinking that the mere operation of the motor centers is insufficient for motor imagery. A bare motor plan is either a dynamic plan (a temporal sequence of motor commands or muscle tensions), or a kinematic plan (a plan for limb movement specified in terms of joint angles). By contrast, motor imagery is a sequence of *faux* proprioception and kinesthesia. *These two things are not the same.* Oddly, this point seems to be underappreciated, so it is worth stressing. In the case of overt movement, motor volleys are sent to the musculature, and this results in movements of the limbs that can be specified dynamically or kinematically. As a functionally distinct part of the movement process, various kinds of sensors, such as stretch receptors and Golgi tendon organs, have the job of measuring aspects of the body in motion. Such measurements result in proprioception and kinesthesia. There is no reason to think that the result of measurement by these sensors produces a signal with even the same number of parameters as the dynamic or kinematic signal, let alone that they are in the same format (i.e., a format such that one can be used as a substitute for the other). The only way to get from the former (signals in motor format) to the latter (signals in proprioceptive and kinesthetic format) is to run the motor signal through something that maps motor plans/signals to proprioception and kinesthesia. And the two possibilities are (a) the body (yielding real proprioception and kinesthesia), and (b) a body emulator (yielding faux proprioception and kinesthesia).

That a motor plan and a sequence of proprioceptive/kinesthetic feelings are distinct should be obvious enough, but the difference can be brought out rather nicely by comparing how the simulation theory and the emulation theory account for phantom limb phenomena. Phantom limb patients fall into two groups: those who can willfully move their phantoms, and those who cannot. While not an entirely hard and fast rule, quite commonly those who cannot move their phantoms suffered a period of pre-amputation paralysis in the traumatized limb, while those who can move their phantoms did not: the trauma resulted in the immediate amputation of the limb with no period of pre-amputation paralysis (Vilayanur Ramachandran, personal communication). Recall the point made in section 2.2 about the requirement that emulators be able to alter their input-output mapping in order to track plant drift. In the case of subjects with a paralyzed limb, the emulator has had a long period where it is being told that the correct input-output

mapping is a degenerate many-to-one mapping that produces the same unchanging proprioceptive/kinesthetic signal regardless of the motor command sent. Eventually, the emulator learns this mapping and the emulated limb becomes “paralyzed” as well. On the other hand, without a period of pre-amputation paralysis, the emulator is never confronted with anything to contradict its prior input-output mapping (to see this point, the difference between [i] a lack of information and [ii] information to the effect that nothing happened, must be kept in mind). On the assumption that phantom limbs are the result of the operation of an emulator, we have a possible explanation for this phenomenon.⁶

Regardless of whether a motor plan is conceived as a dynamic plan or a kinematic plan, it should be clear from the above example that a plan is one thing, and the sequence of proprioceptive/kinesthetic sensations produced when it is executed, is another. The simulation theorist must maintain that those who have an inability to produce motor imagery of a certain sort (because their phantom limb is paralyzed) also have an inability to produce motor plans of that sort. But subjects with paralyzed phantoms seem able to make motor plans – they know all too well what it is that they cannot do; they cannot move the limb from their side to straight in front, for example. What is wrong is that when the plan is executed on the phantom, nothing happens: the proprioceptive/kinesthetic signal remains stubbornly static.⁷ A motor plan is one thing, a sequence of proprioception and kinesthesia is another. The simulation theorist conflates them. This issue will be revisited in the case of visual imagery, where the difference between a motor plan and the sensations resulting from the execution of that plan is clearer.

3.4. Neural substrates

The question of where the neural substrates of the MSS emulators are located is natural. (These emulators are used both in motor control and motor imagery, hence my discussion in this subsection will make reference to both cases.) Masao Ito, who first proposed that emulators would be involved in motor control, postulated that they were in the cerebellum, based on behavioral and lesion data (Ito 1970; 1984). Daniel Wolpert and others (Wolpert et al. 2001) review a range of neurophysiological evidence that also implicates the cerebellum. Motor imagery has long implicated the cerebellum as well as other motor areas. Many such studies are cited in Imamizu et al. (2000), Jeannerod (1995; 2001), and Jeannerod and Frak (1999). In a recent paper that contains useful citations to further neurobiological investigations of MSS emulators, Naito et al. (2002) present results from a range of neuroimaging studies of motor imagery they have conducted, showing that the cerebellum (along with other areas) is selectively active during motor imagery.

4. Visual imagery

So far the applications of the strategy of emulation have been only with regard to motor control and motor imagery. But the same basic scheme is obviously applicable to imagery from other modalities, provided that they have some relevantly exploitable structure. Section 4.1 introduces two models by Bartlett Mel (Mel 1986; 1988), robots that can

use imagery to solve certain problems. Though Mel does not describe them in this way at all, they generate imagery by operating emulators of the motor-visual loop. The details of how these systems generate imagery are no doubt distinct from how nervous systems do it, but at a gross level there is evidence that the same basic architecture is in play. After getting the basic idea across with Mel's models in section 4.1, I will turn in sections 4.2 and 4.3 to providing evidence that some aspects of visual imagery are generated via the operation of a motor-visual emulator in a way at least roughly analogous to that suggested by Mel's models. In section 4.4., I introduce a distinction between modal and amodal imagery. Modal imagery, as exemplified in Mel's models, is imagery based on the operation of an emulator of the sensory system itself, whereas amodal imagery is based on the operation of an emulator of the organism and its environment: something like arrangements of solid objects and surfaces in egocentric space. I show how the two forms of emulation can work in tandem. My goal in these sections is not to provide anything like a systematic or complete theory of visual imagery, but rather to indicate in broad outline how some of the more prominent theories of visual imagery can be seen as implementations of the emulation theory.

4.1. Murphy

Murphy (Mel 1988) is a robot whose job is to move its arm while avoiding contact with obstacles, so that its hand can grasp objects. The arm has three joints – a shoulder, elbow and wrist – all of whose motion is confined to a single plane. There is a video camera trained on the arm and workspace that drives a 64 x 64 grid of units, each effectively a pixel of an image of the workspace. Murphy controls the limb on the basis of the image projected on the grid, where the arm, target, and obstacles are all clearly represented. Murphy operates in two modes. In the first mode, it simply moves its arm around the workspace until it manages to find a path of movement that gets its hand to the target without impacting on any obstacles. Because the arm has redundant degrees of freedom, it is not a trivial problem to find a path to the target. Often what looks initially like a promising route ends up being impossible to manage, and Murphy must backtrack, attempting to work its limb around obstacles in some other way.

The twist is that each unit in the visual grid is actually a connectionist unit that receives an input not only from the video camera, as described, but also receives connections from neighboring units, and a copy of Murphy's motor command (e.g., increase elbow angle, decrease shoulder angle). During overt movement, the units then learn to associate these patterns of input with the future inputs they receive from the video camera – they are learning the *forward mapping* of the motor-visual loop. That is, the grid learns that if the visual input at t_1 is x_1 , and motor command m_1 is issued, the next visual input, at t_2 , will be x_2 . Qualitatively, Murphy's overt motor-visual system is a plant, implementing a forward mapping from current states and motor commands to future states. The visual grid units monitor the inputs to this system (the motor commands) to see what outputs the system produces on their basis (in this case, the system's outputs are patterns of activations on the visual grid).

After a certain amount of experience solving movement problems overtly by trial and error, Murphy gains a new ability, and the use of this ability is the second mode of Murphy's operation. When the visual grid has learned the forward mapping, Murphy is able to solve the problems off-line using visual imagery. It gets an initial visual input of the workspace, including the configuration of the arm and location of target and obstacles. It then takes the real arm and camera off-line, and manipulates the visual grid with efference copies. It moves the image of its arm around by means of the same motor commands that would usually move its arm around, seeing what sequences of movement impact upon objects, sometimes backing up to try another potential solution, until it finds a path that works. At that point, it puts the real arm and camera back on-line and implements, in one go, the solution.

Mel nowhere puts any of this in terms of control theory or forward mappings, et cetera. Rather, he describes it simply as a connectionist network that learns to solve problems through imagery. Nevertheless, during the imagery phase it is clear that the connectionist network is implementing a pseudo-closed-loop control scheme. The grid itself actually serves double duty as both the medium of real visual input and the emulator of the motor-visual loop. When operating on-line, the grid is driven by the video camera. When off-line, it is driven by the activity of its own units and the motor inputs. Because the grid is used for both, the system has a capacity that it never in fact uses. Specifically, it never operates in anything like a Kalman filter mode. This would involve having the imagery capacity engaged during overt operation. In such a mode of operation, the grid would always form an expectation of what the next visual input would be on the basis of the current visual representation and the current motor command. This expectation would then take the form of some degree of activity in some of the units anticipating activation from the camera. This would be helpful in cases where the video input was degraded, and forming an intelligent anticipation would be crucial in interpreting the input.

The next model, also by Mel, is similar to Murphy.⁸ It consists of a robot with two video cameras (to provide ocular disparity), each of which drives an input grid similar to Murphy's. This robot has no limbs, but rather moves itself around wire-frame objects. For example, it might move towards or away from a wire-frame cube, or circle around it. And, just as with Murphy, there are two modes of operation. There is an initial overt mode, during which the robot moves around various wire-frame objects. All the time, the units of the visual grid are getting activation not only from the video cameras, but also from efference copies and from connections to other units in both grids. Again, the grids learn the forward mapping of the motor-visual loop by learning to associate current visual inputs plus motor commands with future visual inputs. Once this is complete, the robot is able to engage in visual imagery in which it can mentally rotate, zoom, and pan images, including images of novel shapes. Upon receiving an initial image on both visual grids from some object, the system takes its motor system and video cameras off-line, and drives the visual grid with efference copies. It can mentally rotate the image by issuing a command that would normally circle the object. It can zoom into or out of the image by issuing a command that would (overtly) move towards or away from the object.

As with Murphy, Mel does not couch any of this in terms of control loops or emulators. And again, the potential for exploiting the grids as part of a Kalman filter-like mechanism for processing perception is not explored.⁹

4.2. Visual imagery and visual areas of the brain

In the KF-control scheme, the emulator is a system that processes sensory information. Specifically, it produces an expectation, and combines it with the sensory residual in order to yield a best estimate of the state of the observed process. For now, the details of the perceptual situation are not the focus. Rather, the point is merely that the emulator is involved both in imagery and in perceptual processing. Mel's models are concrete examples of systems in which the emulator does double-duty, even though in Mel's models the emulators never do both simultaneously.

The hypothesis that this scheme is used by real nervous systems makes a number of predictions, the first of which is that visual "perceptual" areas will be active during visual imagery. And indeed, there is much evidence not only that such areas are active, but that their activity is selectively similar to the activity of such areas during the analogous overt perceptual situations. Because the focus is currently on imagery that is modality specific (see sect. 3.4), the relevant visual areas will include early visual areas.

A number of researchers have reported finding activity in primary visual areas during visual imagery (Chen et al. 1998; for a recent review, see Behrmann 2000; see also Kosslyn et al. 1993). Kosslyn et al. (1995) found that visual imagery not only activates primary visual cortex, but that imagining large objects activates more of this area than does imagining smaller objects, indicating that it is not only active as during imagery, but that details about the kind of activity it presents are also parallel.

In an extremely suggestive study, Martha Farah (Farah et al. 1992) reports on a subject who was to have a unilateral occipital lobectomy. The subject was given a number of imagery tasks before the operation, including tasks in which she was asked to imagine moving towards objects (such as the side of a bus) until she was so close that the ends of those objects were at the edge of her imagined "visual field." After the removal of one of her occipital lobes, the subject was re-tested, and it was found that the distance increased. This suggests that, much as in the case of Mel's models, the image is actually produced on a topographically organized medium, and manipulated via efference copies. With a smaller screen, "walking towards" an imagined object reaches the point where the edges of the object are at the edges of the topographic medium at a greater distance than with a larger screen.

4.3. Imagery and motor control

As Mel's models suggest, some kinds of visual imagery might, surprisingly, require the covert operation of motor areas. In this section I will point out some evidence indicating that, in fact, motor areas are active during, and crucial to, certain sorts of visual imagery. Activity in premotor areas has been widely shown to occur during visual imagery tasks requiring image rotation. Richter et al. (2000) demonstrated this with time-resolved fMRI, a result confirmed by Lamm et al. (2001).

Though such studies are interesting, the theory I am articulating here makes predictions more detailed than the simple prediction that motor areas are active during visual imagery. It makes the *specific* prediction that they are active in producing motor commands of the sort that would lead to the overt counterpart of the imagined event. Enter a set of experiments done by Mark Wexler (Wexler et al. 1998) in which subjects were engaged in imagery tasks while simultaneously producing an overt motor command.

In these experiments, subjects had to solve problems already known to involve certain kinds of visual imagery, specifically the mental rotation of visually presented shapes. At the same time, subjects were to hold and apply a torque (twisting force) to a handle. Results showed that when the required direction of image rotation and the actual applied torque were in the same direction, performance on the imagery task was much better than in trials in which the directions were different. On the emulation theory of imagery, the result is expected. Part of what would be required in order to emulate an event of rotating a shape is a covert motor command such as twisting with one's hand. This is presumably more difficult to do when the motor systems are actually producing a different command than when they are actually producing the required command.¹⁰

A check written at the very end of section 3.3 can now be cashed. As mentioned there, a motor command is one thing, and the character of the sensory signal it produces is something else. Even if this were not clear in the case of motor imagery, it should be quite clear in the case of visual imagery. While motor areas are involved in imagery as described above, clearly the motor plan *by itself* cannot determine the nature of the imagery. Presumably imagining twisting a "d" and a "b" involve identical motor plans – twisting the grasping hand left or right. But the nature of the image produced is quite different – as it would have to be to solve the problem. The difference in this case is that *the states of the emulators in the two cases are different*, and so, driving them with the same motor command does not yield the same result. One yields a rotated "d," the other a rotated "b." Exactly this underdetermination is present in the case of motor imagery, though it is more difficult to recognize here because the body is relatively stable (which is why cases involving plant drift, paralysis of phantom limbs, etc., are good ways to bring the point out).

4.4. Modal imagery versus amodal imagery

An emulator is an entity that mimics the input-output function of some other system. But even when the same control loop is involved, different systems might be being emulated. In Mel's model, for example, the elements emulated are the pixels on the visual input grid, and the relevant dynamics concern the way in which one pattern of active pixels plus a motor command leads to a new pattern of active pixels. Nowhere does the emulator have a component that corresponds to the arm, or hand, or elbow angle. A given active pixel might correspond to part of the hand at one time, and an obstacle at another time. This is, of course, an entirely legitimate use of emulation to produce imagery, in this case specifically visual imagery. The visual input grid is a system of states, and there are rules governing the transitions from one state to the next.

But the emulation might also have taken a different form.

For example, it might have taken the form of an emulator with components corresponding to parameters of the arm itself. This system would learn how these parameters change as a function of the motor command sent to them – hence the forward mapping learned would not be (current pixel pattern + current motor command) → (next pixel pattern), but rather (current arm parameters + current motor command) → (next arm parameters). The overall system would then subject the arm emulator’s state to a “measurement” that would capture the way in which the real video camera maps arm states to grid images.

How can two *different* systems be legitimate objects of emulation like this? As I mentioned, the visual grid is a set of elements whose states change over time in predictable ways. Given visual grid pattern p_1 , and motor command m_1 , the next visual input pattern will be p_2 . This is a forward mapping that is relatively stable, and hence can be learned and emulated. But behind the scenes of this visual grid system is another system, the system consisting of the arm, workspace, video camera, and so forth. This system also consists of entities whose states interact and change over time according to laws – in this case the laws of mechanics and dynamics. And as such, it too implements a forward mapping that can be learned and emulated. And it is obviously true that the visual grid system has the rules of evolution that it has because of the nature of the arm/workspace system and its laws of evolution. If the arm were heavier, or made of rubber rather than steel, then there would be a different mapping from visual input grid patterns plus motor commands to future visual input patterns. Which system is being emulated – the topographic visual input grid or the arm and workspace – is determined by the number and nature of the state variables making up the emulator, and the laws governing their evolution. Mel’s Murphy, for example, uses an emulator whose states (levels of activation of pixel

units) obviously correspond to the elements of the visual input grid.

Either way, the end result is a system that can produce visual imagery. But they do it in different ways. The Mel type systems produce it by constructing an emulator of the sensory input grid itself, and they do this by letting the same hardware act as both the emulator and the sensory topographic input grid. In such a case, the emulator’s state is just the visual image, there is no measurement; or if you like, the measurement consists of the identity matrix. The other potential system I described produces visual imagery by constructing and maintaining an emulator of the *arm’s* state and the *workspace’s* state, and then subjecting this to a “visual measurement,” similar to the measurement that the video camera subjects the real robotic arm to, in order to produce a mock visual image.

Both ways are shown in Figure 7, in which three control loops are represented. The top, boxed in the dashed-line box, is just the actual system. The process is the organism and its environment. The process changes state as a function both of its own dynamic as well as motor commands issued by the organism. The organism’s sense organs produce a measurement of the state of this, resulting, in the visual case, in a topographic image on the retina or primary visual cortex. Nothing new here.

The second, in the box with the dot-and-dash line, corresponds to a modality-specific emulator, as exemplified in Mel’s models. This emulator’s states are just states corresponding to elements in the topographic image. So long as the elements in the image itself compose a system whose states are to some degree predictable on the basis of previous states and motor commands, it implements a forward mapping that can be emulated. Given that the emulated system is just the visual input medium, no measurement is needed in order to arrive at a visual image.

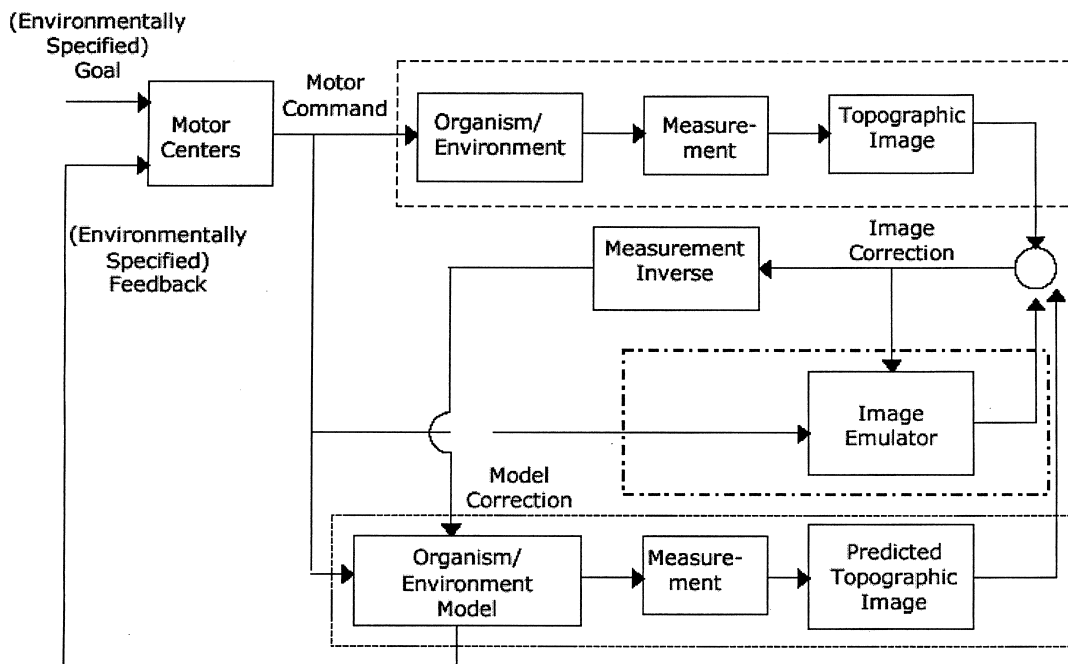


Figure 7. A KF-control scheme using two emulators: one a modality-specific “image” emulator; the other an amodal organism/environment emulator.

The third, in the dotted-line box at the bottom, represents an emulator of the organism and its environment, and for simplicity we can assume that this consists of a sort of model of solid objects of various sizes and shapes in the organism's egocentric space (this will be discussed more in sect. 5).

To run through an example, suppose that we have an organism with a visual sensory modality standing directly in front of a white cube, and it moves to the right. We can describe what happens in the top dashed-line box by pointing out that a white cube is two meters in front of the organism, which is causing a square pattern of stimulation on the retina. The creature moves to the right, and so the cube is now slightly to the creature's left (this is the change in the "organism/environment" box), and the new retinal image (produced by the "measurement") is a square pattern of stimulation on the retina, but slightly shifted to the right from where it was before the movement.

The green-boxed material represents a Mel-type sensory emulator. This system consists of a grid of pixels corresponding to the visual input. Initially it has a square pattern of activation. It processes a "move right" efference copy by altering the pattern of active pixels according to something like the following rules: A pixel p is active at t_2 if and only if the pixel to its immediate left was active at t_1 . This will have the effect of sliding the image one pixel to the right. Given that the emulator in this case consists of nothing but the topographic image, no measurement is needed (or if you like, the measurement consists of the identity matrix). If this is being operated off-line, the resulting image is just visual imagery. If it is on-line, the resulting image is an a priori estimate of the "observed" visual input, which will be combined with the sensory residual to yield the final state estimate.

What about the bottom dotted-line box? Here there is an inner model that represents the organism's egocentric environment: in this case, a cube directly in front of the organism. Subjecting this to a "visual measurement" yields a square topographic image. The "move right" efference copy alters the state of the model, so that the object is now represented as being in front and slightly to the left of the organism (this is the change in the "organism/environment model"). Subjecting this model to a "visual measurement" yields a new topographic input image similar to the previous one, only with the patterns altered slightly. When operated off-line, the result would be visual imagery, if measured; or amodal imagery, if not subjected to a measurement. When operated on-line, the state of this model would constitute the a priori estimate of the layout of the organism's environment, to be modified to some extent by the sensory residual, if any, to yield the final state estimate of the organism's environment.

The two methods are not incompatible, and we could easily imagine a system that uses both, as in Figure 7. This system would run two emulators, one of the sensory periphery as in Mel's models, and also an emulator of the organism/environment, as described above. This system would have not one but two a priori estimates, which could be combined with each other and the sensory residual in order to update both emulators.

An amodal emulator (in this example, the organism/environment model) supplies a number of advantages, stemming from the fact that its operation is not (completely) tied to the contingencies of the sensory modality. First, the or-

ganism/environment model can easily represent states that are not manifest to the sensory periphery. For example, it can easily represent an object as moving from the front, around to the left, behind, and then up from the right hand side as the organism turns around, or objects behind opaque surfaces. This is not something that could easily be done with a Mel-type modality-specific system. In a system that includes both a modal and amodal emulator, the amodal emulator could provide an anticipation in such cases, such as that the leading edge of a square will appear in the visual periphery as the spinning continues.

Second, the same amodal emulator model might be used with more than one modality-specific emulator. I will not bother with a diagram, which would be rather cumbersome, but the idea is that an organism that has, say, visual and auditory modality-specific emulators might be able to run both in tandem with an amodal emulator. In such a case, the amodal emulator would be subject to two different "measurements": a visual measurement, yielding an expectation of what should be *seen*, given the current estimate of the state of the environment; and an auditory measurement yielding an expectation of what should be *heard*, given the current estimate of the state of the environment. And the amodal emulator would be updated by both sensory residuals, resulting in a state estimate that effectively integrates information from all modalities as well as a priori estimates of the state of the environment (Alain et al. 2001; van Pabst & Krekel 1993).

There are additional possibilities that could be explored here. For now I just want to point out that the scheme I am articulating here allows for (at least) two very different kinds of imagery: modality-specific imagery, which is the result of running an emulator of the sensory modality off-line (as in Mel's models), and amodal imagery, which results in the off-line operation of an emulator of the behavioral environment without a corresponding modality-specific measurement. Such amodal imagery might be accompanied with modality-specific imagery, but it might not. More will be said about such cases in sections 5 and 6.¹¹

4.5. Discussion

There are a number of aspects of visual imagery that have not been covered in the discussion of this section. For example, I have said nothing about how a system decides when using imagery is appropriate. I have not mentioned anything about how the imagination process gets started – in Mel's models they begin with an initial sensory input that is subsequently manipulated via imagery, but clearly we can engage in imagery without needing it to be seeded via overt sensation each time. Furthermore, many sorts of visual imagery do not obviously involve any sort of motor component, as when one simply imagines a vase of flowers.

An emulator by itself does not decide when it gets operated on-line versus off-line. Presumably there is some executive process that makes use of emulators, sometimes for imagery, sometimes for perceptual purposes. I am not making the outrageous claim that the brain is *nothing but* a big emulator system. Of course other processes are required. But they would be required in any other account of imagery and perception as well.

Now, once it has been decided that the emulator should be run off-line, it is presumably up to some other system to seed the emulator appropriately. Again, this is a process

outside the scope of the present focus. The initial state of the emulator gets seeded somehow, perhaps with a memory of a state it was once in. An emulation architecture is necessarily part of a larger cognitive system that includes memory and executive processes. My theory is that when this cognitive system is engaged in imagery, it is exploiting an emulator, and when perceiving (see next section), it is using emulators as part of a KF-control scheme. The fact that there are connections to broader cognitive components is not a weakness of my account, but rather a necessary feature of any account of imagery and perception. Detailing such connections would be one of many tasks required in order to completely fill out the account I am outlining in this article.

In the theory I am pushing, visual imagery is “mock” input generated from the operation of an internal emulator. The imagery thus produced depends on what sequence of states the emulator goes through. And this depends on at least three factors. The first, which I will mention and drop, is whatever executive process there is that can initially configure the emulator, as mentioned in the previous paragraph. The second is the emulator’s own internal dynamic; depending on what is being emulated, the state might or might not evolve over time on its own in some specific way. The third factor is efference copies of motor commands. In this section I have focused on imagery produced through the emulation of processes that have no (or a minimal) dynamic of their own, but depend for their evolution on being driven by efference copies. Mel’s models highlight this, as does the imagery involved in Wexler’s studies. But this has been merely a focus of my discussion, not a necessary feature of the model. Some bouts of imagery might involve configuring the emulator to emulate a static process, such as looking at a vase of flowers, where nothing changes over time. In this case, there would be neither any emulator-internal nor efference copy-driven dynamic to the emulator’s state. It would be constant, and yield the more-or-less constant mock sensory input of a vase of flowers. In other cases, there might be a dynamic driven by the emulator’s state, as when I imagine pool balls hitting each other. In this case, the imagined scene is dynamic, but the dynamic is not driven by any efference copies, but by the modeling of processes that evolve over time on their own (Kosslyn [1994] refers to this sort of dynamism in imagery as “motion encoded,” as opposed to “motion added,” which is driven by efference copies). The model thus includes these other sorts of imagery as degenerate cases. I focus on the case involving efference copies to bring out the nature of the fullest form of the model.

Furthermore, the ability of the scheme to handle both modal and amodal imagery allows for explanations of various imagery phenomena. Some sorts of imagery are more purely visual than spatial, as when you simply imagine the colors and shapes of a vase of flowers. Such imagery need not involve imagining the vase of flowers as being anywhere in particular, and might be something like the operation of a purely modal, visual emulator, perhaps exploiting early visual areas. There is a difference between this sort of case and a case where you imagine a vase of flowers sitting on the desk in front of you. In this case, the imagined vase not only has its own colors and textures, but it is located in ego-centric space – you might decide where a vase should be placed so as to obscure a picture on the desk on the basis of such imagery, for example. This might involve both the

modal and amodal emulators, perhaps early visual areas, as well as systems in the dorsal stream. And some tasks requiring spatial imagery might not involve any notably visual component at all.¹² Differing intuitions about whether or not imagery is involved in this or that case might be the result of thinking of different kinds of imagery, kinds that can all be described and explained in the current framework.

A final comment before moving on. The present theory offers a single framework within which both motor and visual imagery can be understood. This is remarkable in itself, because, surprisingly, the dominant accounts of motor and visual imagery in the literature are not at all parallel. Dominant views on motor imagery, as we have seen, equate it with the covert operation of *efferent* processes – either an efference copy or a motor plan. Dominant views on visual imagery treat it as the covert “top-down” stimulation of *afferent* areas. That these two dominant explanations seem so *prima facie* different is at best surprising, at worst embarrassing. First, intuitively, both kinds of imagery seem to be cases of a similar kind of process, and this is no doubt why they are both called “imagery.” Second, a unified account is more theoretically parsimonious, showing how different phenomena can be explained by a single framework. Third, and most important, there are empirically confirmed similarities between imagery in different modalities, such as isochrony (for more such similarities and details, see e.g., Jeannerod & Frak 1999, and references therein). If motor and visual imagery are subserved by different kinds of mechanisms, then such parallels have no obvious explanation. To the extent that it can be shown that similar information-processing structures underlie imagery regardless of modality, then commonalities such as isochrony have a natural explanation.

The theory defended here unifies both accounts seamlessly. Imagery is the result of the off-line operation of emulators. In the motor case, such emulators are predominantly driven by efference copies, especially all cases of motor imagery studied where subjects are invariably asked to imagine *engaging in some motor activity*. Hence the ubiquitous involvement of motor areas in motor imagery. Visual imagery also involves the off-line operation of an emulator (in this case, a motor-visual emulator). But in *some* cases the motor aspect is minimal or absent, because the emulation required to support the imagery does not require efference copies (as when one imagines a vase of flowers), though of course in other cases it does (as in Wexler’s studies). The mechanisms in both cases are the same. The difference is that the kinds of tasks emphasized in visual imagery often lack motor components, whereas motor imagery tasks nearly always highlight exactly the efference-copy driven aspects of the mechanism.

5. Perception

5.1. Sensation versus perception

Psychologists and philosophers have often distinguished between sensation and perception. The distinction is not easy to define rigorously, but the general idea is clear enough. Sensation is raw sensory input, while perception is a representation of how things are in the environment based upon, or suggested by, this input. So, for example, when looking at a wire-frame cube, the *sensory* input consists of twelve co-planar line segments: four horizontal, four

vertical, and four diagonal, arranged in the familiar way. What one *perceives* is a cube, a three-dimensional object in space. That the perception is an interpretation of the sensory input is highlighted by the fact that one can, at least in some cases, switch which face of the cube is in front, as with the Necker cube. Here there are two different interpretations that can be placed on the same sensory input; two different perceptual states based on the same sensory state.

The sorts of representational states that result from perception are extremely complex, but for purposes of the present discussion I will focus on what I take to be the core aspects. Through perception we become aware of objects in our surroundings. A bit more specifically, we become aware of some number of objects and surfaces, their rough sizes and shapes, their dynamical properties (especially movements), and their egocentric locations. To have some handy terminology, I will refer to this as an *environment emulator*. Clearly one of the primary functions of perception is the formation of an accurate estimate of the environment, and this will be embodied in the environment emulator.

Look again at Figure 7. In section 4, I highlighted one aspect of this diagram – its combination of modal and amodal emulators. But now I want to draw attention to another aspect, which is that the feedback from the emulator to the controller does not go through the measurement process. In Figure 2, the control context within which we started involved a controller that was given a goal state, and got feedback that was used to assess the success of the motor program in achieving that goal state. In the feedback control scheme, the feedback is necessarily whatever signal is produced by the plant's sensors, and this imposes a requirement that the goal specification given to the controller be in the same format as the feedback, for only if this is the case can an assessment between the desired and actual state of the plant be made. That is, the goal state specification had to be in sensory terms.

In the pseudo-closed-loop scheme of Figure 4, and the KF-control scheme of Figure 6, the idea that the feedback sent from the emulator to the controller was also in this “sensory” format was retained. In the latter case this was made explicit by including a “measurement” of the emulator's state parallel to the measurement of the real process in order to produce a signal in the same format as the real signal from the plant.

But retaining this “measurement” is neither necessary nor, in many cases, desirable. The real process/plant has many state variables, only a small sampling of which are actually measured. In the biological case, access to the body's and environment's states through sensation is limited by the contingencies of the physiology of the sensors. A system with an amodal emulator that is maintaining an optimal estimate of all the body's or environment's relevant states is needlessly throwing away a great deal of information by using only the mock “sensory” signal that can be had by subjecting this emulator to a modality-specific measurement. There is no need to do this. The emulator is a neural system: any and all of its relevant states can be directly tapped.¹³ This is the meaning of the fact that in Figure 7 the feedback to the controller comes directly from the emulator, without the modality-specific “measurement” being made.

The practical difference between the two cases is significant, because, as already mentioned, a modality-specific measurement process might very well throw out a great

deal of useful information. But the conceptual difference is more important for present purposes. It is not inaccurate to describe the “measured” or “modal” control schemes, including the KF-control scheme of Figure 6, as systems that control *sensation*. Their goal is a sensory goal, they want their sensory input to be like thus-and-so, and they send out control signals that manage to alter their sensory input until it is like thus-and-so. The information they are getting is exclusively information about the state of the sensors. But in the unmeasured amodal variant, the controller has its goal specified in terms of objects and states in the environment, and the feedback it gets is information about the objects in its egocentric environment.

The less sophisticated systems are engaged with their sensors. This is true both on the efferent and afferent ends. In the emulation theory, such systems are engaged in *sensation*. The more sophisticated systems have their goals set in terms of objects and locations in the environment, and receive information in terms of objects and locations in their environment. In the emulation theory, these systems are engaged in *perception*.

5.2. The environment emulator

If the relevant emulator for perception were an emulator of the sensory surface, as in Mel's models, then there would be little question concerning their states – they are just the states of the components of the sensory organs, just as the units in Mel's simulations are pixels of a visual image. But I have claimed that perception involves the maintenance of an emulator whose states correspond to states of the emulated system. This can be made sense of readily in the case of proprioception and kinesthesia of the MSS, as done above in section 3. The relevant states are the dynamic variables of the MSS. But what about other sorts of perception, such as visual perception? What is the emulated system, and what are its states, if not the sensor sheets? To a plausible first approximation the emulated system is the organism and its immediate environment, specifically, objects and surfaces of various sizes, shapes, and egocentrically specified locations, entering into force-dynamic interactions with each other and the organism. Of course, the proposal that an organism's brain has the wherewithal to maintain models with these characteristics does not imply that such models are always complete or accurate. Indeed, there is reason to believe that such models are often schematic, incomplete, and inaccurate. But for my present purposes, the degree of completeness is not important. So long as it is agreed that brains are not always without any kind of model of their environment (and surely this much is unobjectionable), the question concerning what such models are and how they are implemented is unavoidable. How detailed, complete, and accurate such models are in what kinds of circumstances is a crucial issue, but not the current issue.

This involves a combination of *where*, *what* and *which* systems. The *what* and *where* systems are posited to be located in the ventral and dorsal processing streams respectively (Ungerleider & Haxby 1994). The ventral *what* stream proceeds from early visual areas in the direction of the temporal lobes, and appears to be concerned with identifying the *type* of object(s) in the visual field and their properties. The dorsal *where* stream proceeds from early visual areas to the parietal areas, and is primarily concerned with

the location of objects. In addition to these, a *which* system, whose task is indexing and tracking object identity through changes in locations and properties, also appears to be in play (Pylyshyn 2001; Yantis 1992).¹⁴

These systems comprise the core of the environment emulator in the present account. During perception they jointly maintain an estimate of the relevant layout of the environment, especially the number, kind, and egocentric locations of objects. Anticipations of how this layout will change are continually produced, both on the basis of the organism's own movements that result in changes in the egocentric location of objects, as well as anticipated changes brought about by the dynamics of the objects themselves (hence identifying the kind of object in question is crucial). This estimate provides a framework for interpreting sensory input, and is subject to modification on the basis of sensory information.

5.3. On perception and imagery

In contrast to traditional theories of perception that treat it as largely a bottom-up process driven entirely by the sensory periphery, Stephen Kosslyn has developed one of the most influential accounts of the nature of visual perception and its relation to visual imagery. Kosslyn's theory is complex and has many features I will not mention. But I want to focus for now on the main factor that, according to Kosslyn, accounts for the "top-down" nature of visual perception: visual imagery. Kosslyn maintains that imagery processes are used to aid perceptual processing, by filling in missing information on the basis of expectations, for example. As Kosslyn and Sussman (1995) put it, the view is "that imagery is used to complete fragmented perceptual inputs, to match shape during object recognition, to prime the perceptual system when one expects to see a specific object, and to prime the perceptual system to encode the results of specific movements." In fact, not only does Kosslyn claim that visual imagery is an ingredient in visual perception, but he also includes what he calls "motion added" imagery (Kosslyn 1994), which is imagery whose character is in part determined by off-line motor commands.

Although Kosslyn does not cast any of his points in terms of the emulation framework or Kalman filters, it should be clear that the emulation theory is particularly well-suited to explaining the information-processing infrastructure of Kosslyn's account. In the first place, the emulation framework explains exactly how it is that the same process that can operate off-line as imagery, can operate on-line as the provider of expectations that fill in and interpret bare sensory information. In this case, it is the a priori estimate that is then corrected by the sensory signal. Furthermore, the emulation framework's explanation of imagery as the driving of emulators by efferent motor areas in fact predicts what Kosslyn's theory posits, that in many cases the imagery that is used to aid in perception is the product, in part, of the activity of motor areas.

Striking confirmation of this view of the relation between motor-initiated anticipations and perception comes from a phenomenon first hypothesized by von Helmholtz (1910), and discussed and verified experimentally by Ernst Mach (1896). Subjects whose eyes are prevented from moving and who are presented with a stimulus that would normally trigger a saccade (such as a flash of light in the periphery of the visual field) report seeing the entire visual scene mo-

mentarily shift in the direction opposite of the stimulus. Such cases are very plausibly described as those in which the perceptual system is producing a prediction – an a priori estimate – of what the next visual scene will be on the basis of the current visual scene and the current motor command. Normally such a prediction would provide to specific areas of the visual system a head start for processing incoming information by priming them for the likely locations of edges, surfaces, et cetera. Just less than one hundred years after Mach published his experimental result, Duhamel, Colby, and Goldberg (Duhamel et al. 1992) published findings to the effect that there are neurons in the parietal cortex of the monkey that remap their retinal receptive fields in such a way as to anticipate immanent stimulation as a function of saccade efference copies. That is, given the current retinal image and the current saccade motor command, these neurons in the parietal cortex anticipate what the incoming retinal input will be. In particular, a cell that *will* be stimulated because the result of the saccade will bring a stimulus into its receptive field, begins firing in anticipation of the input, presumably on the basis of an efference copy of the motor command. (This is a neural implementation of the construction of an a priori estimate.)

While Kosslyn's theory of visual perception is very detailed, I hope that it is at least clear how these main aspects of his theory can be synthesized by the emulation framework into a yet broader account of the brain's information-processing structure, and, in particular, how the emulation framework makes clear the relation between Kosslyn's theory and specific results in cognitive neuroscience such as those of Wexler et al. (1998) and Duhamel et al. (1992).

5.4. Discussion

In a sense, what I have said about perception glosses almost entirely over what most researchers take to be most important. A standard and unobjectionable view of what perception involves is that it is the creation of a representation of the layout of the organism's environment from bare sensation. In my account, a good deal of this is represented in the KF-control diagram by the line that goes from the sensory residual to the a posteriori correction. The key here is the "measurement inverse," which is just a process that takes as input sensory information, and provides as output information in terms of the states of the environment or the environment emulator. In the case of the amodal environment emulator, this process goes from sensory signals to information about the layout of the environment. This is the process that is the paradigmatic perceptual process, and I say next to nothing about it, except to locate it in the broader framework.

But to fixate on this is to miss the import of the emulation theory. The point is to show how *this* narrow process is part of a larger process, and to do so in such a way as to hopefully highlight two related points. The first is the large-scale nature of perception; the second is the fact that perception is one aspect of a complicated process that intimately involves motor control and imagery. I will address these in reverse order.

The minimal view of perception says nothing at all about how, or even if, perception has any connection at all to systems involved in motor control, imagery, or cognition, and in fact few of the proposals one finds concerning the mechanisms of perception draw any such connections (though,

as Kosslyn shows, some prominent theories do draw such connections). The present account, by contrast, argues that the brain engages in a certain very flexible and powerful sort of information-processing strategy, one that simultaneously addresses all of these (for a Kalman filter model of visual processing compatible with the emulation theory, see Rao & Ballard [1999]). Surely, to treat perception, imagery, and motor control as functionally distinct modules, as though any of them could do its job without the others, is to significantly distort the genuine neurophysiological phenomena.

This leads to the second point, which is that the current scheme, exactly because it treats perception as one aspect of an integrated information-processing strategy, sheds light on the nature of perception itself. In the first place, the scheme highlights the extent to which the outcome of the perceptual process, the state estimate embodied in the emulator, is tuned to sensorimotor requirements. The emulator represents objects and the environment *as things engaged with* in certain ways as opposed to how they are considered apart from their role in the organism's environmental engagements. The perceived environment is the environment as made manifest through the organism's engagements, because the emulator that supplies the perceptual interpretation is an emulator of the agent/environment interactions. The conceptual significance of this is that it allows us to acknowledge the action/behavioral bias of perception without becoming anti-representationalists about perception.

Another shift in emphasis suggested by this account is that perception is shown to be not a matter of starting with materials provided in sensation and filling in blanks until a completed percept is available. Rather, completed percepts of the environment are the starting point, in that the emulator always has a potentially self-contained environment emulator estimate up and running. This self-contained estimate is operational not only during imagery, but presumably also during dreaming (see Llinas & Pare 1991). The role played by sensation is to constrain the configuration and evolution of this representation. In motto form, *perception is a controlled hallucination process*.¹⁵

6. General discussion and conclusion

In this final section, I will make some *very rough* suggestions as to how the emulation theory of imagery and perception not only sheds new light on a number of issues in these fields, but also how it might synthesize aspects of other domains, such as reasoning and language, as well.

6.1. Perception and imagery

The imagery debate, well-known to cognitive neuroscientists, is a debate concerning the sort of representations used to solve certain kinds of tasks. The two formats typically under consideration are *propositions* and *images*. As is often the case, definitions are difficult, but the rough idea is easy enough. Propositions are conceived primarily on analogy with sentences, and images on analogy with pictures. In its clearest form, a proposition is a structured representation, with structural elements corresponding to singular terms (the content of which prototypically concerns objects) and predicates (the content of which prototypically concerns

properties and relations), as well as others. This structure permits logical relations such as entailment to obtain between representations. On a caricature of the pro-proposition view, perception is a matter of turning input at the sensory transducers into structured language-like representations; cognition is a matter of manipulating such structured representations to draw conclusions in accord with laws of inference and probability.

By contrast, images are understood as something like a picture: a pseudo-sensory presentation similar to what one would enjoy while perceiving the depicted event or process. Perception is a matter of the production of such images. Cognition is a matter of manipulating them.

According to the present theory, one of the central forms of imagery is amodal spatial imagery. It will often be the case that this imagery is accompanied by modality-specific imagery, for the same efference copies will drive both the modality-specific emulators, as well as the amodal spatial emulator. Indeed, the fact that there are in-principle isolatable (see Farah et al. 1988) aspects to this imagery may not be introspectively apparent, thus yielding the potentially false intuition that "imagery" is univocal.

Amodal spatial imagery is not a clear case of "imagery" as understood by either the pro-proposition or pro-imagery camps; nor is it clear that such representations are best conceived as propositions. Like propositions, this imagery is structured, consisting at least of *objects* with *properties*, standing in spatial and dynamical relations to each other (Schwartz 1999). They are constructs compositionally derived from components that can be combined and recombined in systematic ways. An element in the model is an object with certain properties, such as location and motion, and this is analogous in some respects to a proposition typically thought of as the ascription of a property to some object.

Nevertheless, such imagery is emphatically *unlike* a picture. This is difficult to appreciate largely because we typically, automatically and unconsciously, interpret pictures as having spatial/object import. But, strictly speaking, this import is not part of the picture. A picture or image, whether on a television screen, a piece of paper, or topographically organized early visual cortex, consists of a dimensionally organized placement of qualities, such as a pattern of colored pixels on a CRT. Seeing such a pattern as representing moving objects, such as Olympic sprinters in a race, involves interpreting the image. The picture itself has no runners, only pixels. Similarly, bare modal imagery is unstructured, lacks any object/spatial import. But because of the potentially close ties between modal and amodal imagery, modal imagery is typically, automatically and unconsciously, given an interpretation in terms provided by amodal imagery.

On the other hand, amodal spatial imagery *is* a representation of the same format as that whose formation constitutes perception, for the simple reason that perception just is, in my account, sensation given an interpretation in terms of the amodal environment emulator. Therefore, although amodal imagery is not picture-like, it is not obviously sentential or propositional either. These amodal environment emulators are closely tied to the organism's sensorimotor engagement with its environment. The model is driven by efference copies, and transformations from one representational state to another follow the laws of the dynamics of movement and engagement (see Schwartz 1999), not of logic and entailment (as typically understood), or at

least not *only* according to logic and entailment. Unlike a set of sentences or propositions, the amodal environment emulator is spatially (and temporally) organized.

I do not have any answers here. I mean merely to point out that if in fact amodal object/space imagery is a core form of neurocognitive representation as the emulation theory suggests, then this might go some way to explaining why two camps, one insisting on understanding representation in terms of logically structured propositions, and the other in terms of picture-like images, could find themselves in such a pickle. The camps would be trapped by the two dominant metaphors for representations we have: pictures and sentences. I am suggesting that neither of these metaphors does a very good job of capturing the distinctive character of amodal imagery as understood by the emulation theory, and that if progress is to be made, we might need to abandon these two relic metaphors and explore some new options, one of which I am providing.

6.2. Cognition

Kenneth Craik (1943) argued that cognition was a matter of the operation of small scale models of reality represented neurally in order to anticipate the consequences of actions, and more generally to evaluate counterfactuals. Phillip Johnson-Laird (1983) has refined and developed this approach under the title of *Mental Models*, and it is currently a dominant theory of cognition in cognitive science. Johnson-Laird describes mental models as representations of “spatial relations, events and processes, and the operations of complex systems,” and hypothesizes that they “might originally have evolved as the ultimate output of perceptual processes” (Johnson-Laird 2001). The representations embodied in the amodal environment emulators are of exactly this sort.

Johnson-Laird’s mental models, while arguably based on something like the representations made available through amodal emulators, involve more than I have so far introduced. Specifically, on Johnson-Laird’s account they are manipulated by a system capable of drawing deductive and inductive inferences from them. The difference between a mental-models account and an account that takes reasoning to be a matter of the manipulation of sentential representations according to rules of deduction and probability is thus not that logical relations are not involved, but rather that the sort of representation over which they operate is not sentential, but a spatial/object model. Exactly what is involved in a system capable of manipulating models of this sort such as to yield inferences is not anything that I care to speculate on now. I merely want to point out that the individual mental models themselves, as Johnson-Laird understands them, appear to be amodal space/object emulators, as understood in the current framework. If this is so, then the emulation framework might be a part of our eventual understanding of the relation between cognition and sensorimotor, perceptual, and imagery processes.

In a similar vein, Lawrence Barsalou (Barsalou 1999; Barsalou et al. 1999) has tried to show that what he calls “simulators” are capable of supporting the sort of conceptual capacities taken to be the hallmark of cognition. Barsalou’s simulators are capacities for imagistic simulation derived from perceptual experience. He argues that once learned, these simulators can be recombined to produce “simulations” of various scenarios, and that such simula-

tions not only subserve cognition, but serve as the semantic import of linguistic expressions.

I am not here specifically endorsing either Johnson-Laird’s or Barsalou’s accounts, though I do think that they are largely compatible with the emulation framework, and each has a lot going for it. My point is merely to gesture in the direction in which the basic sort of representational capacities I have argued for in this article can be extended to account for core cognitive abilities.

6.3. Other applications

There are a great number of other potential applications of the emulation framework in the cognitive and behavioral sciences. In a spirit of bold speculation, I will close by outlining a few of these.

6.3.1. Damasio’s theory of reason and emotion. Antonio Damasio (1994) has argued that skill in practical decision making depends on emotional and ultimately visceral feedback concerning the consequences of possible actions. The idea is that through experience with various actions and the emotionally charged consequences that actually follow upon them, an association is learned such that we tend to avoid actions that are associated with negative emotions or visceral reactions. The relevant part of his theory is that it posits an “as-if loop,” implemented in the amygdala and hypothalamus, that learns to mimic the responses of the actual viscera in order to provide “mock” emotional and visceral feedback to contemplated actions. Though Damasio does not couch it in control theoretic terms, he is positing a visceral emulator, whose function is to provide mock emotional/visceral input – emotional imagery.

A payoff of understanding Damasio’s proposal as a special case of the emulation framework in action is that it allows us to take Damasio’s theory further than he takes it himself. If he is right that the brain employs a visceral/emotional emulator, then it is not only true that it can be used off-line, as he describes. It might also be used on-line as part of a scheme for *emotional* perceptual processing. That is, just as perception of objects in the environment is hypothesized to involve a content-rich emulator-provided expectation that is corrected by sensation, so, too, might emotional perception involve a rich framework of expectations provided by the emulator and corrected by actual visceral input. And just as in environmental perception the nature of the states perceived is typically much richer and more complex than, and hence underdetermined by, anything provided in mere sensation, so, too, the emotional emulator might be the seat of emotional learning and refinement, providing the ever-maturing framework within which raw visceral reactions are interpreted to yield the richer range of emotional perception that we gain as we age.

6.3.2. Theory of mind phenomena. Robert Gordon (Gordon 1986) has been the primary champion of the “simulation theory” in the “theory of mind” debate in developmental psychology. The phenomenon concerns the development of children’s ability to represent others as representing the world, and acting on the basis of their representations (Flavell 1999; Wellman 1990). The canonical example involves a puppet, Maxi, who hides a chocolate bar in location A, and then leaves. While out, another character moves the chocolate bar to location B. When Maxi returns, children

are asked where Maxi will look for the bar. Children characteristically pass from a stage at which they answer that Maxi will look at B, to a stage where they realize that Maxi will look at A, because that is where Maxi thinks it is. According to the simulation theory, we understand others' actions in this and similar situations by simulating them; roughly, putting ourselves in their situation and ascertaining what we would do. Such a simulation might well involve placing ourselves in another's perceptual situation (i.e., creating an emulated surrogate environment situation), and perhaps in their emotional situation with something like the emotion emulator discussed in the previous paragraph.

6.3.3. Situated robotics. Lynn Stein (1994) developed a robot, MetaToto, that uses a spatial emulator to aid in navigation. The robot itself was a reactive system based on Brooks' subsumption architecture (Brooks 1986; 1991). But in addition to merely moving around in this reactive way, MetaToto has the ability to engage its reactive apparatus with a spatial emulator of its environment to allow it to navigate more efficiently. By building up this map while exploring, MetaToto can then use this map off-line (in a manner similar to Mel's models), and can also use it on-line to recognize its location, plan routes to previously visited landmarks, and so forth.

6.3.4. Language. Applications to language are to be found primarily in the small but growing subfield of linguistics known as cognitive linguistics. The core idea is that linguistic competence is largely a matter of pairings of *form* and *meaning*; *form* is typically understood to mean phonological entities, perhaps schematic, and *meaning* is typically understood to be primarily a matter of the construction of representations similar to those enjoyed during perceptual engagement with an environment, especially objects, their spatial relations, force-dynamic properties, and perhaps social aspects, as well. What sets this movement apart is a denial of any autonomous syntactic representation, and the notion that the semantics is based on the construction of representations more closely tied to perception than propositions.

Gilles Fauconnier (1985) has developed a theory of quantification, including scope and anaphoric phenomena, based on what he calls "mental spaces," which, at the very least, are analogous to spatial/object representations posited here. Ronald Langacker's Cognitive Grammar framework (Langacker 1987; 1990; 1991; 1999) is a detailed examination of a breathtaking range of linguistic phenomena, including quantification (the account builds on Fauconnier's), nominal compounds, "WH," passive constructions, and many dozens more. Karen van Hoek (1995; 1997) has developed a very detailed account of pronominal anaphora within Langacker's Cognitive Grammar framework. Leonard Talmy (Talmy 2000), Lawrence Barsalou (Barsalou 1999; Barsalou et al. 1999), George Lakoff, and Mark Johnson (Lakoff 1987; Johnson 1987; Lakoff & Johnson 1999) have also produced a good deal of important work in this area, all of it arguing forcefully that the semantic import of linguistic expressions consists in representations whose structure mimics, because derived from, representational structures whose first home is behavior and perception – exactly the sorts of representational structures made available by the various emulators described here.

6.4. Conclusion

The account I have outlined here is more schematic than I would ideally like. Ideally there would be both more detail at each stage, and more evidence available in support of such details. In some cases such details and evidence have been omitted for reasons of space; in other cases the details and evidence are not currently extant. The primary goal, however, has been to *introduce and articulate* a framework capable of synthesizing a number of results and theories in the areas of motor control, imagery, and perception, and perhaps even cognition and language, rather than to provide compelling data for its adoption. This synthesis is useful for a number of reasons. In the case of motor control, imagery, and perception, many researchers have assumed connections between these phenomena, but have not yet had the benefit of a single framework that details exactly how they are connected. And in the domains I have touched on even more superficially in sections 6.1–6.3, many researchers also make frequent appeal to their models being based in, or being continuous with, motor control, imagery, and perception, but again have not had the benefit of a synthesizing model that helps to make the nature of such connections perspicuous. While space has prevented going into these applications in detail, I hope that my brief remarks in these sections have at least made it clear in outline form how these "higher" brain functions might be synthesized with these other capacities.

In addition to this synthesizing potential, the emulation theory also manages to extend certain ideas in clearly useful directions, such as providing concrete ways of thinking of modal and amodal imagery and their potential interactions, and allowing more clarity regarding the mechanisms of motor imagery. (Even if one thinks that the simulation theory is correct, it is useful to have apparatus to clearly state it in such a way as to distinguish it from the emulation theory. Perhaps this can even lead to better experiment design.)

These considerations are not theoretically insignificant, but they are also quite far from conclusive, or even, on their own, terribly persuasive. Ultimately, of course, informed and detailed investigation will determine the extent to which this framework has useful application in understanding brain function. To date, motor control is the only area in which this framework has the status of a major or dominant theoretical player that has been robustly tested and largely vindicated. I believe that part of the reason for this is that it is only in this area that theorists are generally familiar with the relevant notions from control theory and signal processing, and hence are thinking in terms of this framework at all when interpreting data or designing experiments. Perhaps when more researchers in a wider range of fields are familiar with the emulation theory and its potential applications, it will receive the kind of experimental attention that would be needed to determine the extent to which it is in fact used by the brain as I have claimed.

ACKNOWLEDGMENTS

I am grateful to the McDonnell Project in Philosophy and the Neurosciences, and the Project's Director Kathleen Akins, for financial support during which this research was conducted. I am also grateful to a number of people for helpful feedback on earlier versions of this paper, including Jon Opie, Adina Roskies-Davidson, Anthony Atkins, Brian Keeley, and Jonathan Cohen. A number of referees for this journal also provided extremely valuable suggestions and feedback that resulted in many improvements.

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 com-

plexities 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-

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 Balasubramanian & 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 Balasubramanian 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 sys-

tem 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




Figure 1 (Campbell & Pettigrew). A Kalman filter-control scheme involving two emulators, one modal and one amodal, which can accurately predict refference. The difference between a modality-specific measurement of the amodal emulator and the output of the modal emulator is the predicted refference.

failure to make a clear distinction between the parallel signal emanating from the motor system, the *corollary discharge*, and the precisely calculated sensory consequences of the motor act, the *effference copy* (see Fig. 1).

Whenever a motor command is sent out, a *corollary discharge* – what Grush refers to as an *effference copy* – is generated. This corollary discharge is sent to the emulators so that they may update their representations. From this corollary discharge it is possible to predict the refference that will be generated. *Refference* is a term that refers to any self-generated sensory events that do not correspond to true changes in the world. An inverse of this predicted refference is fed into the sensory pathways, where it cancels out the refferent portion of the sensory input (reviewed in Grusser 1995). The predicted refference is known as the *effference copy*. It is this process that leads to the perceived stability of the world during eye and body movements.

How the brain predicts refference is a fascinating problem that spans motor control (Miall & Wolpert 1996), sensory physiology (Adler et al. 1981; Blakemore et al. 2001; Haarmeier et al. 1997), neuroethology (Bell et al. 1983; 1997), and even psychiatry, where it is postulated that a perturbation of the CDEC system may underlie the positive symptoms of schizophrenia (Blakemore et al. 2002; Feinberg 1978). It has been postulated that one way to solve this problem would be to have an emulator that specifically modelled the expected refference of every motor command and then compared the output with the sensory inflow. Any residue left after the cancellation is *ex-afference* and thus corresponds to true changes in the environment (Blakemore et al. 1998; 2001; Miall & Wolpert 1996).

There are, however, several problems with such theories. As Grush points out (sect. 2.2), in order for emulators to deal with changes in the input/output properties of the body, they must be

modifiable. To be of any use, the emulator must be able to track the changing properties of the body and be able to adjust its properties in step with those of the body. One way to do this is via a Kalman filter (see sect. 2.3). Unlike other emulators, one that solely predicts refference could not utilize a Kalman filter. This is because the output of the emulator cannot simply be compared with the sensory inflow because the prediction is of only a portion of the sensory inflow (the refference). If the emulator is not kept in calibration, it is impossible to tell whether the residue, left after the predicted refference has cancelled with the sensory inflow, is a result of either a real change in the environment or a plant drift. Although such a system may work on a gross scale, such as deciding whether a movement was self- or externally generated, it cannot make fine predictions like those needed by the visual system. It is clear that Grush himself makes a similar mistake to that of Blakemore, Miall, and others when he claims in section 5.2 (last para.) that the output of an environment emulator: “provides a framework for interpreting sensory input, and is subject to modification on the basis of sensory information.” Because an amodal environment emulator will not model refference, there is no way in which it could either help to interpret the sensory input, or be recalibrated by the sensory input – such recalibration would only make the model less accurate.

This problem should not, however, be taken as evidence against Grush’s synthesis, since it clear that it can be extended to predict refference without being affected by the recalibration problems of Blakemore, Miall, and Wolpert. Such an extension would involve comparing the output of two emulators, one modal and one amodal. Under conditions that will generate refference, the output of the amodal and the modal emulator will differ. This is because, whilst the amodal emulator is only predicting the *ex-afference*, the modal emulator is faithfully replicating the sensory

input, *reafference and all*. Thus, it is possible to predict reafference by subtracting the output of the amodal emulator from that of the modal emulator (see Fig. 1). In this schema both emulators can be continually recalibrated via Kalman filters.

3. Neural correlates of Grush's formulation? It is widely accepted amongst neurophysiologists that hindbrain structures such as the cerebellum and torus semicircularis are involved in the calculations of efference copy from the corollary discharges arising as a result of "motor" activity (Bell et al. 1983; 1997; Bodznick et al. 1999; Quiaia et al. 1999). Although Grush does not deny this well-established wisdom, we feel that it would be helpful to the field if he would try to speculate, even at the risk of getting it wrong, about the hindbrain-midbrain-forebrain connectivities that would be required by his schemata, if it is correct that the cerebellum (and hindbrain adnexa) is the site of the efference copy (reafference) calculation. We feel that such an exercise, fraught with potential errors as it may be, could prove a useful antidote to the current philatelic fashion of looking with functional magnetic resonance imaging (fMRI) for some "new" area in the cerebral cortex that is associated with a particular function. It seems unlikely to us that restricting attention to the forebrain will enable us to come to grips with the organizational principles that must underlie the subtle, integrated calculations of ex-afference and reafference considered here.

ACKNOWLEDGMENTS

Our work is supported by the Stanley Foundation. We would also like to thank Guy Wallis.

Duality's hidden influences in models of the mind

Eric Charles

Psychology Department, University of California at Davis, Davis, CA 95616.

epcharles@ucdavis.edu

<http://psychology.ucdavis.edu/grads/Charles>

Abstract: Dualistic approaches to the mind-body relationship are commonplace; however, the adoption of dualistic thinking can often obscure aspects of the way the organism functions as a whole biological entity. Future versions of the emulation theory will, it is hoped, address some of these issues, including the nature of process noise, how distinct iterations can occur, and how to deal with non-emulated aspects of motor control.

The purpose of Grush's target article is to "introduce and articulate" the emulation theory. I wish to mention some obstacles that emulation theory will need to work on which I believe are clouded by the mind-body dualism inherent in the theory. Pointing out some of these seemingly unconnected consequences provides opportunity for clarification and adjustment in future versions of the emulation theory of the mind. I also wish to clarify a point regarding the potential rejection of the emulation theory.

The nature of process noise. In the emulation theory, "process noise" refers to anything causing deviations between actual body position and emulated position. Because body is affected by noise, but mind runs its emulation separately, sources of noise are considered completely unknowable and can only be guessed post hoc when states are compared by the Kalman filter (KF) at a later point in time. This seems unproblematic when one accepts a simplistic mind-body dualism (with its analogous ship/crew relationship), but there are important aspects of the information carried by afferent nerves which it misses. The body continuously sends signals to the mind not only regarding its position and how it got there but also on what affected it. Causes of noise (external effects not anticipated at the start of motion) are often knowable and behave predictably.

Let us say that I reach for a glass on a boat and that an unanticipated wave begins to lift the ship as I start my reach, introducing

noise. Because of the plasticity of my joints, my hand will momentarily be moving too low (as the ship rises the cup rises, but inertia prevents my hand from rising as much). If the KF adjusts for this, it will tell my arm to raise my hand, but as the wave passes, the ship will lower below its initial condition, causing my hand to go well above the cup through over-correction. The emulator is in luck, however, as other parts of the sensory system should send the mind information as to the nature of the noise it is experiencing (from afferents in my feet feeling the upward push of the hull to my inner ear registering the tilt it produces). The wave, although unanticipated, is relatively predictable in effect, and the mind can adjust behavior accordingly and still successfully grab the glass. I believe that most sources of "process noise" could be thus dealt with to a beneficial degree, as large amounts of "noise" are often predictable from minor initial variation. This aspect of control could be missed by thinking of the body as a ship's hull or the mind as a robotic arm control.

Perhaps a combination of the proposed "environment emulator" with the motoric emulator would start to address this issue. Unfortunately, interaction amongst emulators (or whether there are multiple emulators vs. one giant emulator) was not discussed in the target article. Further discrimination of pragmatically different types of noise may also be useful. In either case, I do not believe that would solve the problem completely.

Emulation as an iterative process. The KF performs its functions through a series of discrete "measures" taken from the afferent nerves (both those involved in the motion and those of other senses) when a given time step has past. This is used to estimate the current accuracy of prediction and to attune the filter to the current situation. The KF learns over time, adjusting its gain and calibrating the strength of muscles, flexibility of joints, and so on. The converse of this, of course, is that the body receives signals from the mind in an intermittent manner, adjusting what it is doing only when the mind's KF completes a new comparison, determines the most probable act of correction, and sends the next efferent signal to the body. Clearly, this type of thinking is possible only if the body is one thing and the mind is another thing altogether. But without a suggested mechanism of discrete sampling, one might suspect that feedback must be continuously flowing in both directions, and that both the body and mind were physical components of a single organism. Can the KF model be extended to produce continuous feedback? How would that influence our interpretations of these models or how we build robots/simulations to test them? Also, without requiring a description of a physical instantiation of the KF (see "Cautionary Rejection" below), how would our interpretation of the KF be affected if we viewed it as present in the physical structure of the mind, that is, if we were forced to reject dualism and its accompanying references to mental imagery?

Non-brain mechanisms of adjustment. A final thing that mind-body duality can obscure is neuronal non-brain mechanisms for maintaining and adjusting behaviors. These include feedback loops through the spinal cord, perceptual adjustments to detect sources of noise or gain information which could permit more accurate emulation (such as pupil dilation and analogous auditory adjustments), and so on. There are also a host of non-neuronal bodily adjustments which affect behavior such as hormonal changes and depletion of energy reserves and oxygen supply in local muscular regions. Certainly an emulation could not incorporate many of these, as many of them are activated only by things that would qualify as "process noise" and are therefore unpredictable. Also, many of them act in shorter times than required by the emulation theory (which Grush indicates is approximately 1 second). Are these mechanisms simply not deemed part of motor control? Does simple motor control not require an emulator? An affirmative answer to either question does not seem to mesh well with the emulation model, yet they are clearly implied by the use of dualism.

Cautionary rejection. As a final note, I am concerned with Grush's suggestion (with slight variations in wording in different

places) that “As for what [sorts of] evidence would count against the emulation framework’s applicability to [a domain such as] motor control, it is whatever evidence would count against [the specific model or theory that he is discussing at that time]” (sect. 2.6, para. 3). In these cases, he seems to be using the validity of the particular KF he cites as the only possible validity for an emulation theory of the mind more generally. This is both a brash move on his part (which I applaud) and not particularly historically accurate. It should be clear from the above comments that KFs are at best a highly useful (in the sense of generating novel scientific research) metaphor for how the brain is operating. The target article admits both that there are many different KF models and that necessarily, “the emulation framework relaxes the strict requirements of the Kalman filter” (sect. 2.4, para. 6). Other researchers may support different views of emulation (and if they do not now, certainly they may in the future). These others would certainly not want their opinions rejected if the KF is rejected. More realistically, progressive research programs often modify their predictions. And as Lakatos (1970, p. 151) amply stated: “To give a stern ‘refutable interpretation’ to a fledgling version of a programme is a dangerous methodological cruelty. The first versions may even ‘apply’ only to non-existing ‘ideal’ cases.”

This seems to me the position of a strict KF model of the mind. Certainly it will go through adjustments and changes and, it is hoped, continue to make novel predictions at each stage. I see no reason why future versions of emulation theory could not find solutions to the problems I have pointed out, and I hope that I may look forward seeing them try.

ACKNOWLEDGMENT

Although I have avoided a competitive approach, the criticisms presented here are inspired by James Gibson’s (1966) approach to understanding perceptual systems, which attempts in its own unique way to deal with these issues.

Epistemology, emulators, and extended minds

Terry Dartnall

Department of Computing and Information Technology, Griffith University,
Nathan, Queensland 4111, Australia. terryd@cit.gu.edu.au
<http://www.cit.gu.edu.au/~terryd>

Abstract: Grush’s framework has epistemological implications and explains how it is possible to acquire offline empirical knowledge. It also complements the extended-mind thesis, which says that mind leaks into the world. Grush’s framework suggests that the world leaks into the mind through the offline deployment of emulators that we usually deploy in our experience of the world.

Grush endorses Kosslyn’s claim that when we perceive something in the world we are running an online emulator that fills in information on the basis of expectations. We can then use efference copies of motor commands to run the emulator offline. This continuity between online and offline emulation explains the following epistemological puzzle. You go into a room and see a partially completed jigsaw puzzle on a table. You look at the puzzle and leave the room. You then mentally rotate one of the pieces and discover where it fits. You have now discovered something new – where the piece fits into the puzzle. I think you have discovered it by performing an inner analogue of an operation that, if you had performed it in the world, would have given you an empirical discovery and that also gives you an empirical discovery when you perform it in your mind, even though, in this case, you did not have access to the puzzle.

We can certainly perform such rotations, as R. M. Shepard and associates showed in a series of classic experiments (Cooper & Shepard 1973; Shepard & Metzler 1971). And the operations give

us knowledge: We acquire knowledge by seeing things all the time, most obviously through our straightforward recognition of things in the world. The intuition that we need to overcome is that you derived the knowledge inferentially, *from what you already knew*. So suppose that, rather than leaving the room, you rotate the piece manually and discover where it fits. This is straightforward empirical discovery. From an epistemological point of view, rotating the piece mentally is no different from rotating it manually – in both cases you do not know where the piece fits until you have performed the rotation. Whatever we say about one we will have to say about the other. The physical case is an empirical discovery that is not derived from previous knowledge. Consequently, the mental case is an empirical discovery that is not derived from previous knowledge. What is unusual about it is that you perform it offline, when you do not have access to the puzzle.

I think this can be explained in terms of Grush’s emulator framework. Grush says that perception involves “a content-rich emulator-provided expectation that is corrected by sensation” (sect. 6.3.1). The imagination (by which I mean our ability to form and manipulate images) uses the same emulator to provide similar content, now driven by efference copies of motor commands. In the case of the jigsaw puzzle, we run the emulator online when we rotate the piece manually and we run it offline, using efference copies, when we rotate the piece mentally. When content and copy are veridical, this offline emulator gives us empirical knowledge of the external world.

Now to extended minds. Andy Clark and Dave Chalmers (Clark 2003; Clark & Chalmers 1998) have recently argued that mind extends into the world through the use of “cognitive technology” or “mindware.” It extends through *cognitive processes* when we use pen and paper to work something out, or when we use a computer, or even when we use language, which Clark thinks was the first technology. And it extends when we use physical objects, or even data structures such as encyclopaedias or CD-ROMs, as *external memory stores*, which we can consult “as needs dictate” (the phrase is Clark’s).

Clark’s and Chalmers’ driving intuition is that if something counts as cognitive when it is performed in the head, it should also count as cognitive when it is performed in the world. We now have a natural complementarity, because my epistemological gloss on Grush’s framework says that if a process gives us an empirical discovery when it is performed in the world, it will also give us an empirical discovery when it is performed in the head. This is in keeping with the spirit of the extended-mind thesis, because it erodes the skin-and-skull barrier between mind and world. But we can fill out the framework even more. Clark and Chalmers say that we use objects and data structures in the world as external memory stores. I think there is a complementarity here as well, inasmuch as we have inner analogues of external objects, which we carry around in our heads and consult as needs dictate.

Why do I think we have inner analogues? First, there is the question of symmetry. We perform cognitive actions in the world, and we perform actions in our heads that we would normally perform in the world. We also use the world as an external data store. If the symmetry carries over, we will have inner analogues of external data stores. Next, the problem with using external objects as memory stores is that they are not portable. Inner analogues, which we could carry around in our heads, would free us from this limitation. But more important, there is this: If we perform operations that we would normally perform in the world, on objects that are not present to our senses, then we must have inner analogues of those objects to perform the operations on. Consider the case of the jigsaw piece when we leave the room. We perform an operation in our minds that we would normally perform in the world. We say, loosely speaking, that we rotate the piece in the imagination. But what do we really rotate? The answer has to be: an inner analogue of the piece. As with the external piece, we can consult this inner analogue as needs dictate. We have remembered knowledge about the piece, so sometimes we will retrieve this remembered knowledge. But sometimes we will retrieve *non-*

remembered knowledge – and we will do so *exactly* when we perform operations on inner analogues that we would normally perform in the world.

How common is this? The key question is whether imagination *in general* is an active process. Perception is an active process of saccading and foveating. If the imagination has taken its cue from perception, as the emulator theory suggests, then it would seem that we regularly saccade and foveate onto inner analogues of external objects to acquire empirical knowledge, as needs dictate. When we ask ourselves whether frogs have lips or whether the top of a collie's head is higher than the bottom of a horse's tail, we foveate onto inner images, just as we foveate onto real frogs and real horses and collies. These kinds of inner operations may be more common than we had thought.

Grush's framework shows how it is possible to have offline empirical knowledge. It also complements the extended-mind thesis. If something counts as cognitive when it is performed in the head, it should also count as cognitive when it is performed in the world (mind leaks into the world). But also, if a process gives us an empirical discovery when it is performed in the world, it will also give us an empirical discovery when it is performed in the head (the world leaks into the mind). I think that Grush's emulator framework shows us how this is possible.

Where in the brain does the forward model lurk?

Opher Donchin^a and Amir Raz^b

^aDepartment of Biomedical Engineering, Johns Hopkins University, Baltimore, MD 21205; ^bDepartment of Psychiatry, Division of Child and Adolescent Psychiatry, Columbia University College of Physicians and Surgeons and New York State Psychiatric Institute, New York, NY 10032.
opher@bme.jhu.edu ar2241@columbia.edu
<http://www.bme.jhu.edu/~opher>

Abstract: The general applicability of forward models in brain function has previously been recognized. Grush's contribution centers largely on broadening the extent and scope of forward models. However, in his effort to expand and generalize, important distinctions may have been overlooked. A better grounding in the underlying physiology would have helped to illuminate such valuable differences and similarities.

Despite the length of this piece, Grush's goal is modest: He attempts to show how seemingly disparate fields can be unified under the conceptual construction of the forward model, or emulator. In his conceptual framework, Grush argues that modeling is a common theme in activities that involve fashioning our own behavior, predicting the behavior of others (i.e., theory of mind), or expecting changes in the environment. Grush implies that this general network manifests in converging neurophysiological mechanisms.

Whereas this idea is not entirely novel, it is interesting to compare Grush's presentation with like accounts that were originally raised more than a decade ago with the advent of a cerebellar role in cognitive functions (Ito 1993; Kawato 1997). Those discussions related the idea of emulation to specific anatomical and physiological details, making testable predictions that are fruitful to this day. In contrast, the target article generally avoids a discussion of the underlying mechanisms, leaving the reader unclear as to the practical significance of the emulation theory.

Grush says that, at least for motor control and motor imagery, the forward model is likely implemented by the cerebellum. The target article would have benefited from a review of evidence suggesting that other modeling functions are also cerebellum-dependent (e.g., theory of mind [ToM]). The cerebellum is one of the brain structures consistently abnormal in autism (Courchesne 1997), concomitant with impairment in ToM (Frith 2001). Moreover, the cerebellum has occasionally been implicated in func-

tional magnetic resonance imaging (fMRI) studies pursuing the locus of ToM (e.g., Brunet et al. 2000). On the other hand, ToM is usually associated with the prefrontal cortex or, possibly, the amygdala (e.g., Siegal & Varley 2002), and most neuroimaging studies do not find cerebellar activation (e.g., Castelli et al. 2002). If mechanisms of ToM are cerebellum independent, does it not have implications for Grush's theory? We feel the author should have addressed the physiological literature much more extensively, perhaps at the expense of other points.

By way of an intellectual detour relevant to issues of the forward model and ToM, we point out the view that impairment of the forward model for motor control may be key to inappropriate behavior (e.g., in psychopathology). In the case of delusions of control (e.g., schizophrenia), abnormal behavior may arise because failure of the forward model causes a perceived difference between expected and veridical consequences of motor commands (Frith & Gallagher 2002; Frith et al. 2000). The role that the forward model of one system might play in the behavior of another system seems relevant to Grush's sweeping theory.

While these issues go unaddressed, Grush devotes considerable attention to his emulation theory of motor imagery (previously suggested by Nair et al. 2003 and Berthoz 1996), contrasting it with the seemingly similar simulation theory. His argument for the emulation theory depends on a critical assumption that motor planning is in either kinematic or dynamic coordinates rather than in sensory coordinates. However, Grush does not convincingly support this assumption, and there is some reason to challenge its validity. For example, recent evidence on the effect of eye position on the behavior and physiology of reaching (Batista et al. 1999; Henriques et al. 1998) has been used to argue that reaching is planned in visual coordinates (Batista et al. 1999; Donchin et al. 2003). Moreover, even if we accept Grush's assumption, he does not explore the inevitable subsequent physiological implications. Presumably, motor planning takes place in either primary motor (MI) or premotor areas, and the forward model is to be implemented by the cerebellum. Towards that end, the actual sensory experience should be in either the primary or the secondary somatosensory cortex (SI or SII). However, fMRI studies of motor imagery find activation of MI, premotor areas, and the parietal reach regions (all regions associated with motor planning), but neither SI nor SII display such compelling activations (e.g., Hanakawa et al. 2003; Johnson et al. 2002; Servot et al. 2002;).

Grush also invests in a detailed development of the Kalman filter. The Kalman filter is an important idea in motor control, where a proper mixture of estimation and feedback are necessary for performance, but it is not appropriate in the other systems. In extending the model from the world of motor control, Grush obscures the fundamental idea behind the Kalman filter: The quality of the signals is used to determine the balance between its inputs. A gating, rather than filtering, mechanism would have been more fitting for all of his other examples, and the implementation of gating mechanisms is a different problem from that of filtering.

The difference between a gated and a filtered system affects the characteristics of the required forward model. The Kalman filter theory of motor control would be effectively served by an unarticulated forward model that calculated a rough linear approximation. This forward model needs to be fast, but it does not need to be accurate (Ariff et al. 2002). In contrast, the forward model implied by the emulation/simulation theory of motor imagery is the opposite: It does not need to be any faster than the actual motor-sensory loop of the body (and evidence indicates that it indeed is not faster; Reed 2002a), but it should provide an accurate notion of the sensations that would accompany action (Decety & Jeannerod 1995). We feel that physiological accounts could speak to such differences, and a more rigorous exploration might have made them more obvious to both Grush and his readers.

In sum, like Grush we agree that modeling is an important brain function. However, we believe that Grush's generalized approach may at times blur important distinctions rather than unravel previously unseen commonalities. We feel that had Grush more

closely tied his account to the physiological literature, this shortcoming might have been evaded.

Emulators as sources of hidden cognitive variables

Peter Gärdenfors

Department of Cognitive Science, Lund University, Kungshuset, S-222 22 Lund, Sweden. Peter.Gardenfors@lucs.lu.se
<http://www.lucs.lu.se/People/Peter.Gardenfors/>

Abstract: I focus on the distinction between sensation and perception. Perceptions contain additional information that is useful for interpreting sensations. Following Grush, I propose that emulators can be seen as containing (or creating) hidden variables that generate perceptions from sensations. Such hidden variables could be used to explain further cognitive phenomena, for example, causal reasoning.

I have a great deal of sympathy for Grush's emulator model. Albeit still rather programmatic, it promises a powerful methodology that can generate a multitude of applications in the cognitive sciences.

Grush presents some evidence concerning the neural substrates of the emulators. However, this evidence is based on different kinds of neuroimaging. In my opinion, one should rather be looking for functional units in the brain, described in neurocomputational terms that can be interpreted as some kind of Kalman filter. At a low level, the example from Duhamel et al. (1992) concerning saccade anticipation seems to be such a system. However, the functional units should be searched for at higher levels of cognition as well. What ought to be developed is a way of combining the modeling techniques of artificial neuron nets with the control theoretical principles of Kalman filters (see the volume by Haykin [2001] for some first steps). What is needed, in particular, is an account of how a Kalman filter can *adapt* to the successes or failures of the controlled process.

As used in traditional control theory, Kalman filters operate with a limited number of control variables. In his presentation in section 2.3, Grush presumes that the emulator has the same set of variables as the process to be controlled. Although he notes that this is a special case and mentions that the variables of the emulator may be different from those of the process itself, he never presents alternative versions of the filters.

Now, from the perspective of the evolution of cognition, the distinction between sensation and perception that Grush makes in section 5.1 is of fundamental importance (Gärdenfors 2003; Humphrey 1993). Organisms that have perceptions are, in general, better prepared for what is going to happen in their environment. My proposal is that perceptions are generated by emulators and they function as forward models.

One important property of an emulator is that it does not need to rely exclusively on the signals coming from sense organs; it can also *add on* new types of information that can be useful in emulating. As a matter of fact, Grush (1998) has written about this possibility himself:

The emulator is free to "posit" new variables and supply their values as part of the output. A good adaptive system would posit those variables which helped the controller [. . .] They are variables which are not part of the input the emulator gets from the target system. They may be the actual parameters of the target system, they may not. But what is important is that *the emulator's output may be much richer than the sensory input it receives from the target system.* (emphasis in original)

It does not matter much if the added information has no direct counterpart in the surrounding world as long as the emulation produces the right result, that is, leads to appropriate control signals.

The information provided by these variables is what generates the difference between sensations and perceptions. For example, when the system observes a moving object, its sensations consist

only of the positions of the object, whereas the *forces* that influence the movement of the object are not sensed. However, if the system has been able to extract "force" as a hidden variable and relates this to the sensations via something like Newton's Second Law, then the system would be able to make more efficient and general, if not more accurate, predictions.

In section 2.2, Grush makes the point that emulators must have a certain degree of plasticity. This is not sufficient: A general theory must also account for how an emulator can *learn* to control a system. Supposedly, it slowly adjusts its filter settings (and set of variables) on the basis of some form of reward or punishment feedback from the process to be controlled. This would be analogous to how artificial neuron networks learn. Such a form of learning may pick up higher-order correlations between input and output. These correlations may be expressed by the hidden variables of the emulator.

The hidden variables of the multimodal emulators that Grush discusses in section 6.1, may provide the system (the brain) with cognitive abilities such as object permanence. More generally, one would expect the multimodal emulator to represent the world in an *object-centered* framework, rather than in a viewer-centered one (Marr 1982). As Grush (1998) writes: "[S]pace is a theoretical posit of the nervous system, made in order to render intelligible the multitude of interdependencies between the many motor pathways going out, and the many forms of sensory information coming in. Space is not spoon-fed to the cognizer, but is an achievement." Another speculation is that phenomena related to *categorical perception* are created by the hidden variables of the emulator.

More generally, different kinds of emulators may produce the variables that are used in *causal reasoning*. An interesting finding is that there is a substantial difference between humans and other animal species. As has been shown by Povinelli (2000) and others, monkeys and apes are surprisingly bad at reasoning about physical causes of phenomena. Tomasello (1999, p. 19) gives the following explanation of why monkeys and apes cannot understand causal mechanisms and intentionality in others: "It is just that they do not view the world in terms of intermediate and often hidden 'forces,' the underlying causes and intentional/mental states, that are so important in human thinking."

On the other hand, even very small human children show strong signs of interpreting the world with the aid of hidden forces and other causal variables. Gopnik (1998, p. 104) claims that "other animals primarily understand causality in terms of the effects of their own actions on the world. In contrast, human beings combine that understanding with a view that equates the causal power of their own actions and those of objects independent of them." Apparently, humans have more advanced causal emulators than other animals.

Finally, as Grush mentions in section 6.3.2, another relevant area is our "theory of mind," that is, the ability of humans to emulate (yes, not simulate) the intentions and beliefs of other individuals. An important question for future research then becomes: Why do humans have all these, apparently very successful, emulators for causes and a theory of mind, and why do other species not have them? A research methodology based on emulators and Kalman filters may provide the right basis for tackling these questions.

From semantic analogy to theoretical confusion?

Valérie Gaveau,^a Michel Desmurget,^a and Pierre Baraduc^b

^aSpace and Action, INSERM U53, Bron, 69500, France; ^bINSERM U483, 75005 Paris, France. gaveau@lyon.inserm.fr
desmurget@lyon.inserm.fr pbaraduc@snv.jussieu.fr
<http://www.lyon.inserm.fr/534>

Abstract: We briefly address three issues that might be important to evaluate the validity of the “emulation theory”: (1) Does it really say something new? (2) Are similar processes engaged in action, imagery, and perception? (3) Does a brain amodal emulator exist?

In this nicely written paper, Grush proposes the “emulation theory of representation” as a unifying principle able to synthesize “a wide variety of representational functions of the brain” (target article, Abstract). This attempt to merge heterogeneous models into a single conceptual framework is meritorious. However, based on the following arguments, we feel that this idea remains highly debatable scientifically, although it is seductive intellectually.

A “new” theory? The idea that common emulators might be used by different functions of the brain is not new, as acknowledged in the target article. For instance, a motor theory of perception has long existed in psychology and has been convincingly supported by behavioral studies (e.g., Viviani & Stucchi 1989; 1992). In the same vein, optimal estimation has long been described as a pivot between the motor and sensory domains. Under this concept are grouped statistical methods devised to extract the valuable part of a noisy signal knowing a priori information. In the central nervous system, this information can be a motor command or an a priori belief on the sensor input. The theoretical concept of optimal estimation (and, in particular, Kalman filtering or KF) has been quite convincingly argued to apply to the analysis of sensory signals, whether visual or proprioceptive (Rao & Ballard 1999; Todorov & Jordan 2002; van Beers et al. 1999; 2002; Weiss et al. 2002; Wolpert et al. 1995), thus establishing a clear link between visual perception and motor control at the level of sensory processing.

The main aim of the emulation theory was to extend this link to a much larger variety of processes. However, only remote and conjectural arguments are presented in the target article with respect to this goal. In other words, there is a nice description of several items of evidence in favor of partial links that have been known to exist for a long time, but there is no clear articulation of these partial links into a general model. In this sense, Grush’s theory cannot be considered truly new. Its scientific support reaches the same boundaries as the previous unarticulated theories. The only articulation between these theories lies in conjectural assertions and in the semantic confusion introduced by terms such as “emulation,” “prediction,” and “estimation.” We do not want to seem excessively discourteous, but all that seems to hold at the end of the article might be something like “emulation processes take place in the brain for various functions.” To make his claim more convincing, Grush has failed to address key issues such as: (1) What, besides the word “emulation,” is common between the predictive activities involved in tasks as different as guiding the hand toward a target (motor control), generating a structured sentence (language), or determining where “Maxi will look” (theory of mind)? (2) What could be the nature of the common substrate that is postulated to be involved in those incredibly dissimilar tasks?

Are similar processes engaged in action, imagery, and perception? One of the main claims of Grush’s article is that the “emulator” used for controlling action can be used for imagery. However, as far as motor imagery is concerned, strong interferences have been demonstrated to exist between actual and represented postures (Sirigu & Duhamel 2001). This may suggest an exactly opposite interpretation of the Wexler experiments (Wexler et al. 1998), which are presented as a key support to the emulation theory. How can Grush rule out the possibility that the conflict takes

place between the *sensory* outcome predicted by the actual motor command (through the forward model) and the mentally rotated one, and not between the actual motor command and the command necessary to rotate the object? It is quite difficult to see how “emulating” the rotation would simply be *possible* when the motor cortex is engaged in a task incompatible with the mental rotation (does this imply the existence of dual forward models?). In contrast, it is understandable that the voluntarily imagined visual scene can dominate the (involuntarily) predicted one, since both are constructs which are unrelated to the actual, static, visual feedback.

In parallel to the previous remarks, we argue that Grush’s model also meets a problem when faced with neuropsychological evidence. Abnormal timing of imagined movements has been found in parietal patients with normal overt movements (Danckert et al. 2002; Sirigu et al. 1996), showing that motor imagery necessitates more than the simple prediction of sensory feedback used for on-line control. A similar conclusion was reached by Schwobel et al. (2002). A functional dissociation seems also to exist between perception and action (Milner & Goodale 1996). For example, movement guidance relying on forward modeling has been shown to be dramatically impaired in a patient presenting with a bilateral posterior parietal lesion (Grea et al. 2002; Pisella et al. 2000). When submitted to standard neurological tests, this patient does not present cognitive or perceptual problems. Trying to extend the concept of emulation to other sorts of imagery is still more problematic. Indeed, dissociations between intact visual imagery and profoundly affected visual perception have been found in several patients (Bartolomeo et al. 1997; Beschin et al. 2000; Goldenberg et al. 1995; Servos et al. 1995). These results openly contradict the notion that visual imagery emerges via an “emulation” of normal vision through top-down processes.

The brain amodal emulator in perception. We were truly puzzled by the suggestion that an *amodal* emulator of the external world could exist in the brain. This claim seems to negate the rich literature documenting dissociations between our different senses (see, e.g., the intermodal conflicts generated by prism adaptation or pinna modification). For instance, biasing the input in a given sensory modality leads to an adaptation of that modality (e.g., the waterfall illusion in vision). It is possible that Grush would interpret this result as a change in the emulator (it is a change in the prior probabilities of object motion, which is part of our knowledge of the world – supposedly analogous to the command of a KF). However, in contrast to the prediction of an amodal emulator, it can be shown that this kind of adaptation does not transfer to other modalities. In fact, besides this remark, what seems to emerge from the recent research is the rooting of high-level supramodal abilities (such as conceptualization) in modality-specific experience (Barsalou et al. 2003).

Does the brain implement the Kalman filter?

Valeri Goussev

Motor Control Laboratory, Rehabilitation Institute, Montreal, H3S 2J4, Canada. valeri@colba.net <http://www.colba.net/~valeri/>

Abstract: The Kalman filtering technique is considered as a part of concurrent data-processing techniques also related to detection, parameter evaluation, and identification. The adaptive properties of the filter are discussed as being related to symmetrical brain structures.

Since the 1960s the data-processing community has been fascinated by the appearance of the new filtering technique (Kalman & Bucy 1961), which had naturally extended Wiener’s filtering theory into the multidimensional time-variant domain. The clarity and simplicity of its structural design has allowed this technique to dominate for more than 40 years in different fields: technology, biology, and economics. Its popularity has grown tremendously.

However, when speaking about data processing in the brain, we should be aware of some essential limitations of the technique which can sometimes be frustrating.

First of all, the Kalman filter is the linear filter. In spite of the fact that there is also experimentally supported evidence of linear transforms, such as the short-interval Fourier transform claimed to be found in the visual cortex (Glezer et al. 1985), there is greater evidence which demonstrates nonlinear data processing in the brain. Even the simple signal operations such as summation and subtraction, commonly used in linear theories, are in doubt when dealing with spike-time sequences that are always positive.

Moreover, even being locally in the linear processing range, we could have a number of other well-known linear tasks: detection, identification, parameter estimation, and pattern recognition. These tasks seem to be significant also for real data processing in the brain, but they cannot be expressed in terms of the Kalman filter, because they use other loss functions. Nevertheless, the target article draws attention to the important particular problem of structural design in filtering technique, which seems to be related in the article more to correct estimation of the plant parameters rather than to the optimal control task itself. The information aspect of the control could take place in the structural organization of the brain.

Following this direction and thereby being further in the scope of the Kalman filtering technique, we could pay attention to other attractive features which seem to have been barely touched upon in the target article. How can one converge to the right filter parameters when the input signal or the mixed noise changes their characteristics? The basic idea underlying the Kalman filtering technique is the lemma of orthogonal projections, which states that for each optimal filter the following equation should be valid:

$$(s - w, z) = 0 \tag{1}$$

where z is the filter input signal, $z = s + n$, s is the useful signal, n is the white noise, w is the filter output signal, and $(*,*)$ is the scalar product.

If the task consists of filtering a low-frequency signal from its additive mixture with the white noise, the equation (1) is equivalent to the requirement for the observed error signal $e = z - w$ to be the white noise too, that is, to have its spectral density constant over all observable frequencies. This property initiated in the early 1960s the appearance of certain specific adaptive filters, converged to the optimal filters (Kalman & Bucy 1961; Wiener 1950). The basic idea of these adaptive controls was the insertion in the

non-optimized Kalman filter of at least two additional band pass filters, the outputs of which were proportional to the spectral density of the filter error at different frequency bands. The difference between two band filter outputs was used as a control signal to adjust parameters of the main filter.

Another, more powerful, optimization technique is based on the structural representation of the gradient for the loss function $l = (e, e) = \min_A$, where $A = \{a_i\}_{i=1}^n$ is the parameter vector of the main filter. The gradient $\vec{G} = \{G_i\}_{i=1}^n$; $G_i = \partial l / \partial a_i$ can be obtained using the model of the main filter (in reference to the target article, we should have two feedback loops including emulators). The model may not be an exact copy of the main filter, but it has the same outputs for parameter disturbances according to the perturbation theory (Bellman 1964). The possible structural design of this optimization technique is presented in Figure 1 for the optimal regulator problem for the infinite-time stationary case, which is also included in the scope of the Kalman filtering technique.

Starting with small interconnection parameters, we have two filters (main filter and its model) that are functionally equal because they deal with approximately equal signals. Gradually increasing parameter values (until 1) leads to inequality in the functional orientation of the filters, giving the main filter the role of transferring the input signal, and the model the role of processing the error. The slightly different optimization structure can be used for the pure optimal filtering problem. Besides the evident analogy with the presence of symmetrical structures in the left and right hemispheres of the brain and their interconnections, there is no other experimental evidence supporting the optimization technique. However, it is difficult to withstand the temptation (following Wiener's famous book; Wiener 1950) to speculate about extending this adaptive mechanism to human society. The more similar are the governing structure (main filter) and the opposition (the model), the more effective and reliable is the optimization process to get to the extremum of the goal criterion.

In conclusion, does the brain implement the Kalman filter? Certainly it should, as far as it deals with filtering tasks. However, one could say that, in spite of the fact that the Kalman filter's basic properties are valid for a number of practical tasks and can also explain some physiological phenomena, their scope is severely limited and not sufficient for understanding even the basics of data processing in the brain. We could expect the appearance of a more general nonlinear theory which will be able to embed the Kalman filter theory, likely as it did with the Wiener filter theory.

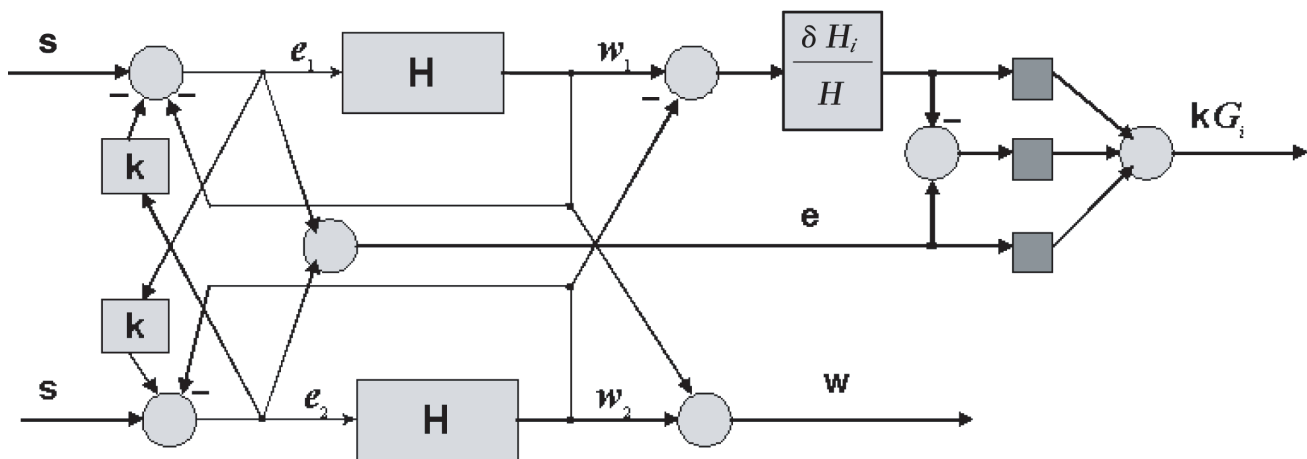


Figure 1 (Goussev). Symmetrical structure to obtain components for the loss function gradient in the optimal regulator problem. H is open loop transfer function, $\partial H_i / H$ is perturbation transfer function for parameter a_p ($\partial H_i = \partial H / \partial a_i$), black squares are power measuring devices.

Amodal imagery in rostral premotor areas

Takashi Hanakawa,^a Manabu Honda,^{b,c} and Mark Hallett^d

^aHuman Brain Research Center, Kyoto University Graduate School of Medicine, Sakyo-ku, Kyoto, 606-8507, Japan; ^bLaboratory of Cerebral Integration, NIPS, Myodaiji, Okazaki, 444-8585, Japan; ^cPRESTO, Japan Science and Technology Agency, Kawaguchi, 332-0012, Japan; ^dHuman Motor Control Section, NINDS, National Institutes for Health, Bethesda, MD 20892-1428. han@kuhp.kyoto-u.ac.jp honda@nips.ac.jp hallettm@ninds.nih.gov <http://hbrc.kuhp.kyoto-u.ac.jp>

Abstract: Inspired by Rick Grush’s emulation theory, we reinterpreted a series of our neuroimaging experiments which were intended to examine the representations of complex movement, modality-specific imagery, and supramodal imagery. The emulation theory can explain motor and cognitive activities observed in cortical motor areas, through the speculation that caudal areas relate to motor-specific imagery and rostral areas embrace an emulator for amodal imagery.

The “simulation” theory of motor imagery has been the primary basis of the idea that movement and motor imagery substantially share the underlying mechanisms and hence the neural correlates. In fact, many neuroimaging studies (e.g., Deiber et al. 1998) have demonstrated that motor imagery evokes activity of efferent motor areas such as the supplementary motor areas (SMA), ventral premotor cortex (PMv), and dorsal premotor cortex (PMd). Each of these motor areas, however, is known to have rostral and caudal subdivisions. It is commonly regarded that caudal cortical motor areas concern relatively simple movements, whereas rostral motor areas control complex movements. Modern anatomical and neurophysiological evidence basically supports this rostral-caudal functional gradient in the motor areas: Caudal zones relate more to the output-oriented components and rostral zones relate more to the preparatory or sensory-cognitive preprocessing components of motor control (Geyer et al. 2000). In a recent study comparing complex motor execution and motor imagery, whereas the caudal PMd and SMA showed movement-predominant activity, the rostral PMd and SMA were active equally for the two conditions, endorsing the significance of rostral motor areas as the controller common to complex motor execution and imagery (Hanakawa et al. 2003c).

A new account for a puzzle. Therefore it was initially puzzling when we found conspicuous activity in the rostral motor areas during various cognitive tasks requiring neither overt motor response nor typical motor imagery (Hanakawa et al. 2002). Moreover, this activity was co-localized with activity during a complex finger-tapping task, which involves the PMd (Sadato et al. 1996). The cognitive tasks used in this study included serial mental addition involving verbal/phonological imagery and also mental operations in two-dimensional space depending on visuospatial imagery. Be-

cause of this diversity, it was difficult to count on the “simulation” theory for the explanation of those motor-area activities. Although the motor-area activities were likely linked to some mental imagery processes, they would not be in a modality-specific form. There were several possibilities to explain those overlapping activities during motor tasks and various sorts of imagery tasks (e.g., general working memory process, arbitrary stimulus-response linkage, etc.), as we have already discussed in Hanakawa et al. (2002).

Now Rick Grush’s emulation theory of representations has put a new account on the top, perhaps, of this list. The emulation theory tells us that a neural system incorporated with an emulator that works for controlling movement and generating motor imagery may be able to produce other sorts of imagery too. Furthermore, the emulation theory allows a specific emulator for “amodal” imagery in addition to the ones for modality-specific imagery. Our results make sense if the activities in the rostral motor areas during those various cognitive tasks represent “amodal” spatial imagery processes, which run parallel with domain specific imagery. By contrast, we can speculate that the caudal motor areas, potentially including the primary motor cortex, are associated with motor specific imagery.

Although it is not clearly indicated in Grush’s paper, modality-specific and nonspecific amodal imagery might be situated in overlapping but slightly different brain regions. There is evidence to support this. The comparison between a serial mental addition task and a number rehearsal task has disclosed that modality-specific imagery (i.e., motor-type imagery associated with silent verbalization) can account for caudal PMd activity observed during both tasks but not rostral PMd activity observed only during the addition task (Hanakawa et al. 2003a). The difference between these two tasks, among other differences, was that the serial addition task demanded more complicated imagery for which subjects needed to operate rigorously on imagery contents. Such complex imagery may overflow to the amodal spatial emulator. In addition, although many people would think that mental calculations should rely exclusively on a phonological-verbal imagery emulator implemented in the linguistic resources, this is not always the case.

Mental abacus and emulation. An interesting situation occurs when experts of abacus operations perform calculations using either a physical or a mental abacus. An abacus is a traditional calculation device consisting of a frame, a horizontal dividing bar, and columns of beads, each of which has a place value. Abacus operators manipulate the beads with their fingers to calculate. That is, numbers are represented as spatial distribution of beads, and the visuomotor control governs the operational rules in the abacus-based calculation.

Intriguingly, abacus experts not only manipulate the device skillfully but also develop an amazing mental calculation skill after proper training. The way of training provides an interesting

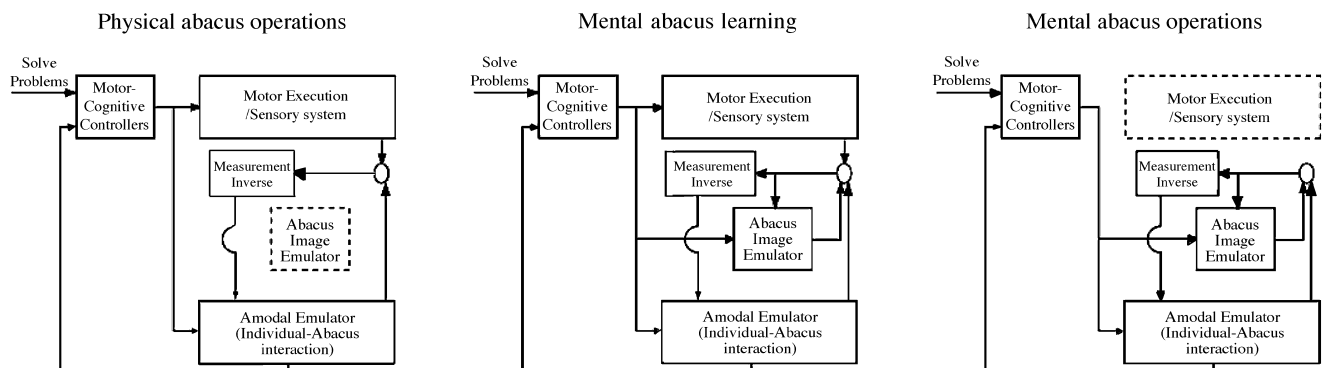


Figure 1 (Hanakawa et al.). An “emulationist” control theory of physical abacus operations, mental abacus learning, and mental abacus operations. The actual motor execution/sensory appreciation system is cut off from the “motor-and-cognitive controllers” during mental abacus operations.

perspective. They first learn physical abacus operations, and then they train themselves to operate on a mental abacus image, moving their fingers as if they were pushing imaginary abacus beads. Once they fully develop the mental calculation skill, they usually do not bother to move their fingers while performing mental calculation. Figure 1 illustrates a control theory of physical and mental abacus operations from the “emulationist” viewpoint adapted from Grush’s Figure 7 in the target article. Based on this theory, mental abacus operations correspond to offline, conscious manipulation of an imaginary abacus supported by a modality-specific emulator. To achieve this, however, amodal imagery is probably working in the background by emulating rules that govern expert abacus interaction and monitoring what is going on in the virtual space.

The neural substrates during mental abacus operations included the rostral PMd, posterior parietal cortex, and the posterior cerebellum, bilaterally (Hanakawa et al. 2003b). Notably, control nonexperts also showed activity in the left rostral PMd and posterior parietal cortex, in addition to the language areas, during mental calculation. This result further supports the amodal nature of imagery computed in the rostral motor areas.

Conclusions. The above-mentioned rostral motor area activities coexist with activity in the posterior parietal cortex and also the cerebellum, to which Rick Grush has tentatively assigned the neural correlates of the “emulator.” Taken together, therefore, rostral motor areas may constitute a part of the neural network representing the “emulators,” particularly of amodal imagery. An alternative explanation for the amodal functions of rostral motor areas may be that these areas correspond to one of the key structures representing the “controllers” for both motor and cognitive operations, as we show in Figure 1.

ACKNOWLEDGMENTS

This work was supported in part by a National Institute of Neurological Disorders and Stroke (NINDS) Intramural Competitive Fellowship Award; and by a Grants-in-Aid for Scientific Research for Young Scientists (B) to Takashi Hanakawa (grant no. 15700257) and a Priority Area grant to Manabu Honda (grant no. 15016113) from the Ministry of Education, Culture, Sports, Science, and Technology of Japan.

The size-weight illusion, emulation, and the cerebellum

Edward M. Hubbard and Vilayanur S. Ramachandran

Department of Psychology, Center for Brain and Cognition, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0109.

edhubbard@psy.ucsd.edu vramacha@ucsd.edu

<http://psy.ucsd.edu/~edhubbard>

<http://psy.ucsd.edu/chip/ramabio.html>

Abstract: In this commentary we discuss a predictive sensorimotor illusion, the size-weight illusion, in which the smaller of two objects of equal weight is perceived as heavier. We suggest that Grush’s emulation theory can explain this illusion as a mismatch between predicted and actual sensorimotor feedback, and present preliminary data suggesting that the cerebellum may be critical for implementing the emulator.

If a person compares the weight of a large object with that of a small object of identical physical weight, the latter will feel substantially heavier, even though the person is explicitly asked to compare the weight rather than the density. This effect – the so-called “size-weight illusion” – is a striking demonstration of the principle that perception is predictive and does not simply involve a passive response to sensory inputs (Charpentier 1891; Ross & Gregory 1970). Traditionally, it has been suggested that the brain expects the bigger object to be much heavier and sets the muscle

tension accordingly, and so when the larger object is lifted it feels surprisingly light (Ross 1966; Ross & Gregory 1970), indeed lighter than a small object of identical weight. However, recent evidence has shown that the size-weight illusion (SWI) persists despite adaptation of these peripheral lifting movements (Flanagan & Beltzner 2000), suggesting that the source of the illusion may be a central mismatch between the expected and actual sensory feedback. We therefore suggest that the source of the mismatch in the SWI may be an internal sensory prediction, which, after a lifetime of experience, generates an erroneous weight prediction, yielding a sensory residual and the corresponding illusion (Hubbard et al. 2000; in preparation).

One prediction that we made on the basis of this hypothesis is that patients with damage to the cerebellum, which has been implicated in weight perception (see Holmes 1917; 1922), may also show reductions in the SWI, even in the absence of impairments in weight perception. A number of researchers (e.g., Kawato 1990; Wolpert et al. 1995; and Grush in the present target article) have suggested that predicting the sensory consequences of motor actions may be a function of portions of the cerebellum, especially the dentate nucleus. These speculations led us to wonder whether the cerebellum may be involved not only in overt movement, but also in cognitive simulation prior to movement, functioning as a “Grush emulator” (Grush 1995; target article), or forward model (Jordan & Rumelhart 1992). This line of reasoning is also supported by the observation that neurons in the lateral cerebellar cortex (specifically lobules V and VI) respond to the anticipated sensory consequences of an action (Miall 1998).

To test this prediction, we tested six control subjects and seven cerebellar patients. Cerebellar patients of varying etiologies were referred to us by physicians on the basis of neurological assessment. Patients showed typical signs of cerebellar dysfunction including intention tremor, past pointing, and dysdiadochokinesia.

To assess weight discrimination, subjects were presented with a pair of weights differing in weight by 50 grams and were asked to state which of the two cans was heavier. We used both a pair of large cans (300 g, 350 g) and a pair of small cans (150 g, 200 g). Each subject was tested twice. After assessing weight discrimination, we assessed the magnitude of the SWI by asking subjects to determine which of 10 small cans (ranging from 100–275 g) matched the apparent weight of the large 300-gram can. Each subject was tested four times.

The six control subjects showed accurate weight discrimination. Subjects made errors on a total of four out of 24 trials, and no subject made more than one error. However, when asked to match the weight of the large 300-gram can, control subjects showed a clear SWI, matching the large can with a can that weighed substantially less (mean 151.04 g). The magnitude of the illusion is far greater than the minimum difference that can be discriminated – the illusion is not due to an inability to discriminate the weight of the cans.

On the other hand, five of seven cerebellar patients showed a reduction of the SWI, despite intact weight discrimination. The first patient, a middle-aged woman showing acute unilateral cerebellar signs (left hand) caused by secondary tumor metastasis in the brain, showed the most dramatic effect. She was mentally lucid, intelligent, and articulate. She showed cerebellar signs – intention tremor, past pointing, and dysdiadochokinesia – only in the left hand. Her ability to estimate subtle differences in weight was identical in both hands. However, the left hand showed no SWI, whereas the right hand showed the illusion in full strength. She expressed considerable surprise that the two hands were producing different results on the task. There was some recovery from cerebellar signs on the following day, and this time the left hand showed the illusion, but still it was substantially smaller than in the normal, right hand.

The subsequent six patients had bilateral cerebellar damage caused by injury, infarction, and electrocution (one patient). Unlike the first patient, they were seen weeks to months after the on-

set of the lesion. Two of these showed the full SWI, despite florid bilateral cerebellar signs. The other four showed a small reduction of the illusion (mean 175 g in the four patients who showed the effect, with almost 210 g in one of them, compared to the 300 g standard).

These experiments provide some preliminary evidence that the cerebellum may be involved in sensory predictions of overt motor behavior and thereby contributes to the SWI. However, these results are not conclusive. Future research should be conducted with additional cerebellar patients to distinguish between three hypotheses: (1) a specific part of the cerebellum serves as a Grush emulator and this was severely damaged in Patient One but less damaged in the others; (2) Patient One may have had some other metastatic lesion causing the reduction of the SWI, for example, a zone in the basal ganglia or frontal lobes that receives information from the cerebellum rather than in the cerebellum itself; or (3) the loss of the SWI may be seen only acutely (as in Patient One) and, given the cerebellum's remarkable adaptive capacities, may have recovered substantially in the other patients. These findings suggest that the cerebellum may be involved in perceptual and cognitive predictions, functioning as a Grush emulator or forward model for internal simulations before performing certain tasks.

The role of “prespecification” in an embodied cognition

J. Scott Jordan

Department of Psychology, Illinois State University, Normal, IL 61960-4620.
jsjorda@ilstu.edu <http://www.ilstu.edu/~jsjorda>

Abstract: Grush makes extensive use of von Holst and Mittelstaedt's (1950) efference copy hypothesis. Although his embellishment of the model is admirably more sophisticated than that of its progenitors, I argue that it still suffers from the same conceptual limitations as entailed in its original formulation.

Efference-copy models tend to be based on a sensory-motor distinction in which the terms “sensory” and “motor” imply functionally orthogonal halves of an organism (Jordan 2003). This habit has its scientific roots in the Bell-Magendie law – the discovery that the spinal cord entails separate ascending and descending tracts (Boring 1950). It was this neurological fact, along with others, that motivated Pavlov's and Sherrington's reflexologies, as well as von Holst and Mittelstaedt's (1950) control theoretic critique of such stimulus–response (S–R) based approaches. And although the efference-copy hypothesis offered a seemingly workable alternative to S–R approaches, it still entails a commitment to functional orthogonality implied in the terms “sensory” and “motor.”

The problem with such proposed orthogonality is that more and more data indicate the nervous system does not function in this way. And what is more, Grush himself touches upon the most robust data in support of this point as he discusses the role the cerebellum might play in his emulator model. In traditional models of the “motor-control” hierarchy, a desired behavior, expressed in body coordinates, is fed from association cortex to the motor cortex. It is then converted into the actual motor command, that is, the torque to be generated by the muscles. This motor command is then sent to both the musculoskeletal system and the spinocerebellum–magnocellular red nucleus system (SMRN). The SMRN system has access to both the motor command and its immediate sensory effects. The SMRN uses these sources of information to generate what Kawato et al. (1987) refer to as a motor-error signal. Because the cerebro-cerebellar loop is faster (10–20 msec; Eccles 1979) than the cerebro-spinal loop, the use of anticipated motor error, or “virtual feedback” as Clark (1997) refers to it, affords control at much finer time scales than that allowed via the “real” feedback obtained through the cerebro-spinal loop.

Though this description makes it sound appropriate to refer to this hierarchical system as a motor-control system, the model itself challenges such language. Specifically, in addition to inputs from the association cortex, the motor cortex also receives input from the SMRN. This SMRN signal, however, defies definition via terms such as “motor” or “sensory.” It is neither, yet at the same time, both. For all intents and purposes, it is best described as Clark described it: as a “virtual feedback” or a virtual effect. Given that this virtual effect figures into the content of the actual motor command, the motor command also defies definition via terms such as “motor” and “sensory.” For all intents and purposes, it is perhaps best described as a specified virtual effect. Perhaps at the motor-cortex level of the hierarchy, this virtual effect is expressed in terms of anticipated or intended torque, whereas at the association-cortex level of the hierarchy, the virtual effect is expressed in terms of a more distal, environmental consequence. The point is that at all levels within this hierarchy, what are being prespecified (i.e., commanded), detected, and controlled, are effects (i.e., feedbacks) that play themselves out at different spatiotemporal levels for different systems.

Grush himself indirectly addresses this point when he argues that sensation and perception both constitute control systems. They both utilize prespecifications (i.e., “goals”) and control feedback. If this is truly the case, however, it means that control systems control their input (i.e., feedback), not their output (Powers 1973). When this notion of prespecified/controlled input is applied to Grush's account of motor control, a contradiction is generated between the notions of prespecified input and an efferent-motor command, for the efference copy is traditionally modeled as a prespecified motor output. Hershberger (1976) was aware of this contradiction and coined the concept “afference copy” to address the fact that control systems prespecify, monitor, and control inputs (i.e., effects/feedback). Hershberger's notion of “afference copy” makes it clear that all of the prespecifications (i.e., goals, control signals, and efference copies) in the system are prespecifications of effects (i.e., input/feedback). The entire system, therefore, seems more appropriately modeled as an effect-control hierarchy.

In addition to overcoming some of the conceptual problems engendered by the efference-copy hypothesis, the notion of effect-control also provides a means of potentially integrating ecological and representational approaches to perception. Grush's model is firmly entrenched in the representational camp. His model begs representationalism because he begins by conceptually dividing the problem into organisms and environments. Given this dualism, the task becomes one of determining how it is that organisms build models of the environment in their brains in order to get around in the world. This then sets the stage for the introduction of yet another dualism – efference and afference.

The notion of effect control provides a means of avoiding such dualisms, for it begins by recognizing that the common denominator among environments, organisms, brains, and neurons, is regularities. Every aspect of an organism, including its nervous system, can be coherently modeled as an embodiment (Jordan 1998; 2000) or encapsulation (Vandervert 1995) of environmental regularities. The implications of this notion are straightforward. There is no need to divide an organism's nervous system into biological and informational properties. Nervous systems are, by necessity, embedded, embodied regularities. The control dynamics of such systems, therefore, need not be modeled via terms such as “sensory,” “motor,” “afferent,” and/or “efferent.” Such terms are used because of our historical commitment to the input-output orthogonality inherent in the Bell-Magendie law. What control systems do is to prespecify and control effects.

Once such an embodied controller is in place, its own regularities become available for further embodiment. Grush acknowledges this point when he argues that because his emulators are neural systems, any and all of their relevant states can be directly tapped. Tapping into such regularities affords an organism the ability to control effects at increasing spatiotemporal scales. At every point in this phylogenetic bootstrapping process, regulari-

ties and their control are the issue at hand. Seen in this light, Gibson's (1979/1986) notion of "resonance," as opposed to representation, takes on new meaning. An organism's nervous system resonates to environmental regularities because the nervous system itself is an embodiment of those regularities.

Computational ideas developed within the control theory have limited relevance to control processes in living systems

Mark L. Latash^a and Anatol G. Feldman^b

^aDepartment of Kinesiology, Pennsylvania State University, University Park, PA 16802; ^bDepartment of Physiology, Neurological Science Research Center, University of Montreal and Rehabilitation Institute of Montreal, Montreal, Quebec, H3S 2J4, Canada. mLL11@psu.edu
feldman@med.umontreal.ca

Abstract: Exclusively focused on data that are consistent with the proposed ideas, the target article misses an opportunity to review data that are inconsistent with them. Weaknesses of the emulation theory become especially evident when one tries to incorporate physiologically realistic muscle and reflex mechanisms into it. In particular, it fails to resolve the basic posture-movement controversy.

There is little doubt that the central nervous system (CNS) takes advantage of previous experiences and implicit or explicit knowledge of physical properties of the body and the environment to perceive and act in a predictive manner (Bernstein 1967; Lashley 1951). The target article describes the currently dominant view that these processes result from the use by the CNS of internal models that compute, imitate, or emulate the input/output relationships, or their inverses, characterizing the dynamic interactions among elements of the organism and between the organism and the environment. The theory is based on a belief that computational ideas developed in areas of robotics and missile control are relevant to biological systems.

The target article presents no convincing experimental evidence in favor of this view for motor control and kinesthesia. For example, contrary to the conclusion by Wolpert et al. (1995), no emulation, with or without Kalman filters, is needed to account for their observation of errors in the hand-position estimation (Feldman & Latash 1982; Ostry & Feldman 2003). The analysis of balancing a pole on a finger (Mehta & Schaal 2002) strongly suggests coupling between motor processes and visual information; such coupling has been interpreted within different theoretical frameworks without invoking the concept of an emulator (Dijkstra et al. 1994; Warren et al. 2001).

Exclusively focused on data that are consistent with the theory, the target article misses an opportunity to review data that contradict it. Weaknesses of the emulation theory become evident when one tries to incorporate physiologically realistic muscle and reflex mechanisms. In particular, there are powerful neuromuscular mechanisms that generate changes in the muscle activity and forces to resist perturbations that deflect the body or its segments from a current posture (Feldman 1986; Matthews 1959). The emulation theory fails to explain how the body can voluntarily change position without triggering resistance of posture-stabilizing mechanisms; hence, it fails to resolve the most basic posture-movement problem in motor control (von Holst & Mittelstaedt 1950/1973).

Specifically, within the emulation theory, to make a movement certain changes in the muscle activity and torques are computed and implemented. These torques move body segments from the initial posture. This produces resistance generated by the mentioned mechanisms that try to return the segments to their initial position. The system may reinforce the programmed action by generating additional muscle activity to overcome the resistance. After the movement ends, muscles cannot relax without the segments returning to the initial position, a prediction that contra-

dicts empirical data (Sternad 2002; Wachholder & Altenburger 1927/2002). The theory overlooks an existing, empirically based solution to the posture-movement problem (Feldman & Levin 1995; see also below).

The target article disregards a basic idea that, to be physiologically feasible, a theory must carefully integrate properties of the neuromuscular system (Bernstein 1947; 1967). In particular, Bernstein made two points: (1) control levels of the CNS cannot directly program performance variables such as muscle activation levels, muscle forces, joint torques, joint angles, and movement trajectories; and (2) any repetition of a motor task is *always* accompanied by non-repetitive neural patterns at any level of the neural hierarchy ("repetition without repetition"). Compare these conclusions with two statements in the target article: "A bare motor plan is either a dynamic plan (a temporal sequence of motor commands or muscle tensions), or a kinematic plan" (sect. 3.3, para. 2) and "a particular motor command might lead to one MSS [musculo-skeletal system] output at one time, but lead to a slightly different output at some time months or years later" (sect. 2.2, para. 4).

Other well-established physical and physiological principles of motor control have been disregarded in the emulation theory, suggesting that the theory has no chance to succeed despite its present popularity. Experimental support for the statements below can be found in von Holst & Mittelstaedt (1950/1973), Bernstein (1967), Latash (1993), Enoka (1994), Feldman and Levin (1995) and Ostry and Feldman (2003).

The target article makes a correct point that properties of muscles and their reflexes are unpredictable in advance. However, this does not justify the conclusion that an emulator is required to correct for such a "deficiency." Instead of ignoring or predicting and correcting these properties, the controller may take advantage of them to assure stable motor performance in the poorly predictable environment. In particular, by resetting the thresholds of activation of motoneurons, the controller may predetermine where, in angular coordinates, muscles start their activation and manifest their reflex and intrinsic properties. This control process underlies the ability to relax muscles at different joint configurations and to change body posture without provoking resistance of posture-stabilizing mechanisms.

Threshold control (a notion absent in the emulation theory) complies with Bernstein's insight that control levels cannot in principle specify commands that predetermine such performance variables as joint torques, joint angles, or "muscle tensions" (a poorly defined notion in the target article). Threshold control may be efficient even if the behavior of effector structures is not perfectly predictable because it specifies equilibrium states of the neuromotor system. It allows achievement of a desired motor effect by providing different neural commands to muscles via reflexes without additional interference from the controller (cf. "repetition without repetition"). The idea of movement production by shifts in equilibrium states avoids many of the problems that emerge when control schemes are borrowed from an area of technology (e.g., robotics) where movements are powered by predictable actuators, not variable, spring-like muscles. Such control does not need an on-line emulator to lead to stable behavior.

The last point is related to the area of motor illusions and imagery. We appreciate the desire of the target article's author to offer a control scheme that accounts for many aspects of human behavior. However, we do not see how such a scheme can predict illusions of anatomically impossible joint configurations or changes in the body anatomy such as those observed under muscle vibration (Craske 1977). One famous example is the so-called Pinocchio effect, when the vibration of the biceps of an arm whose fingertip lightly touches the tip of the nose leads in some persons to a perception of elongation of the nose.

Internal models and spatial orientation

Daniel M. Merfeld

Jenks Vestibular Physiology Laboratory, Massachusetts Eye and Ear Infirmary, Boston, MA 02114. dan_merfeld@meei.harvard.edu
<http://www.jvl.meei.harvard.edu/jvpl>

Abstract: Aspects of “emulation theory” have been seminal to our understanding of spatial orientation for more than 50 years. Sometimes called internal models, existing implementations include both traditional observers and optimal observers (Kalman filters). This theoretical approach has been quite successful in helping understand and explain spatial orientation – successful enough that experiments have been guided by model predictions.

First, a correction is warranted. Typically, the feedback signal from a Kalman filter to a controller is not the estimated noise-free measurement, $I^*(t)$, as shown in Figure 6 of the target article. Because one has the freedom to choose the state variables that are most interesting for a specific control problem, it is standard to choose the state variables $r(t)$ and $r^*(t)$, such that they include the primary states that one is interested in controlling. Therefore, these are generally used for such feedback control. For example, one could choose to make position (e.g., eye position) a state variable for control problems that require accurate control of position (e.g., saccades to a stationary target) but could also choose to make velocity (e.g., eye velocity) a state variable for those problems that require control of velocity (e.g., smooth pursuit or vestibulo-ocular reflexes). One seldom has the same freedom to choose sensors. For example, the semicircular canals cannot directly indicate angular position instead of angular velocity.

More importantly, while agreeing that motor control has been well served by frameworks related to Grush’s emulation theory, I vehemently dispute the notion that “motor control is the only area in which this framework has the status of a major or dominant theoretical player” (sect. 6.4, Conclusion, para. 3). Kalman filtering, observers, and other closely related concepts have been hypothesized to contribute to aspects of human spatial orientation (i.e., balance control, vestibulo-ocular reflexes, and perception of motion and orientation) for more than 50 years. There is no question that these “internal models” have contributed substantially to our understanding of spatial orientation and are major theoretical players on this stage.

Because space considerations prevent comprehensive coverage, I will briefly introduce some models used to understand human spatial orientation. The original authors can be allowed to speak for themselves via extensive citations. Von Holst (1954) proposed a concept, springing from spatial orientation experiments performed in collaboration with Mittelstaedt (von Holst & Mittelstaedt 1950), that has come to be known as “efference copy.” The concept was that efferent commands are copied within the central nervous system (CNS) and used to help distinguish re-afference – afferent activity due to internal stimuli (e.g., self-motion) – from ex-afference – afferent activity due to external stimuli. A similar hypothesis was proposed by Sperry at about the same time (Sperry 1950); some credit von Uexkull (1926) for publishing a similar hypothesis almost 30 years earlier.

Later, while studying adaptation to sensory rearrangement (Hein & Held 1961), Held (1961) recognized a shortcoming in von Holst’s original formulation – that efference copy and the re-afference generated by self-motion cannot be compared directly because one signal has the dimensionality of a motor command whereas the other has dimensions of sensory feedback. Held proposed that efference copy must proceed to a hypothetical neural structure called correlation storage, with the output from the correlation storage compared to the sensory afferent. (The similarity between Held’s “correlation storage” and the process model in Grush’s Figs. 5 and 6 is striking, as is the similarity between Held’s comparator and Grush’s sensory residual difference calculation.)

Reason (1977; 1978) noted that earlier theories had looked at

sensory conflict as an incompatible difference between different sensory afferents and emphasized that sensory conflict is more accurately defined when sensory inputs differ from the expected sensory pattern that we expect, based upon our previous exposure history. These concepts were later included in Oman’s thoroughly considered observer theory model of motion sickness (1982; 1990; 1991). Because Kalman filters, discussed in detail by Grush, are simply observers (Luenberger 1971) whose feedback gains are optimally chosen based on noise characteristics, Oman’s observer formulation and that presented by Grush are indistinguishable.

In parallel with these efforts, Young and some of his colleagues were developing and implementing models of human spatial orientation. Following a study of the interaction of rotational motion cues from the semicircular canals and vision, Zacharias and Young (1981) modeled their responses with the inclusion of an internal model of the canals. This model was later followed by a model (Borah et al. 1988) based on optimal observer theory (Kalman filtering) that was developed to explain and model how the nervous system combined and processed sensory information from several sensory systems (vision, semicircular canals, and otolith organs).

More recently, we (Merfeld 1995a; 1995b; Merfeld et al. 1993; Merfeld & Zupan 2002) have implemented a series of models that most closely follow in the trail blazed by Oman’s observer-theory approach to explain both vestibulo-ocular reflexes and perceived motion and orientation. Another model (Glasauer 1992), making similar predictions and addressing similar questions and problems, was developed in parallel with the above approach. At almost the exact same time, a third group developed a similar approach using what they referred to as coherence constraints (Droulez & Cornilleau-Peres 1993; Droulez & Darlot 1989; Zupan et al. 1994). This approach used internal models, not unlike those used by Merfeld and Glasauer, and might be accurately characterized as a thoroughly considered implementation of Held’s correlation storage and comparator discussed previously. More recently, the observer and coherence constraint approaches have been combined in a comprehensive model of sensory integration (Zupan et al. 2002). In parallel with all of these efforts, a model has been developed to explain and model human postural control using optimal observers (Kuo 1995). This postural model implements internal models of both sensory and motor systems.

Given this lengthy and illustrious history, including implementations of both traditional observers and optimal observers (Kalman filters), it is our opinion that the lack of general acceptance is *not* due to the lack of a formal theoretical or modeling framework, but rather, on account of *too little cross-disciplinary communication* (such cross-disciplinary communication is fostered by *BBS*) and *the lack of direct experimental evidence and experimental investigations specifically targeted to prove or disprove the hypothesized “internal models.”* To address the latter problem we have recently begun performing experimental investigations (Merfeld et al. 1999; Zupan et al. 2000; Zupan & Merfeld 2003) specifically designed to prove or disprove aspects of the internal model hypotheses.

ACKNOWLEDGMENT

The author acknowledges support from NIH/NIDCD grant DC 041536.

The art of representation: Support for an enactive approach

Natika Newton

Department of Philosophy, Nassau Hall, Nassau Community College,
Garden City, NY 11733. natika@worldnet.att.net

Abstract: Grush makes an important contribution to a promising way of viewing mental representation: as a component activity in sensorimotor processes. Grush shows that there need be no entities in our heads that would count as representations, but that, nevertheless, the process of representation can be defined so as to include both natural and artificial (e.g., linguistic or pictorial) representing.

Grush helps banish the misleading practice of viewing “representations” as static entities, and of using “representation” as an object term like “picture” or “imitation,” as in “That picture is not an original Rembrandt, it’s an imitation.” Grush’s theory shows that the term refers to a process. Nothing is “a representation” except as a part of that process, and even then the usage is ambiguous. Think of representation as the process of representing, and of a particular case as an instance of representation; but never think of that instance as “a representation.” I predict that if we are successful at reconceiving representation so that the latter usage becomes technically meaningless, then representation will no longer be a philosophical problem.

Humans are skilled at the art of representation, socially and aesthetically, and traditionally their products are called “representations.” (Grush himself employs that usage at times.) People who are successful create an object that produces in the perceiver a state similar in desired ways to that of perceiving what the artist is representing. In itself, there is no harm in calling the products “representations”; the harm comes when the source of these objects is ignored, and one looks for some feature of the objects *themselves* that makes them representations. Physical resemblance is one obvious candidate. The most dangerous view, for purposes of philosophical clarity, is holding that a natural mapping from the “representations” to the worldly objects they represent exists through causal relations between representations and objects. In natural language, words “represent” objects through the relation between the objects and the learning of the words. Language has become the paradigm for representation for many philosophers, to the extent that many theorists derive the nature of mental representations from that of natural language. There are well-known problems with this approach to representation and no generally accepted solution; this path leads to a dead end.

But representation is a real phenomenon. If it is a process, then we should seek the paradigm in human activity. I (Newton 1996), along with many others (Grush mentions some, e.g., Johnson-Laird 1983; Lakoff 1987), hold that representation is based in human behavior and perception. The general theory is simple to state in terms of motor control systems. An alternative to pure feed-forward or feedback systems is one in which the final goal state is *emulated* by a mechanism allowing the system to perform its goal-directed actions in a type of *rehearsal* of the decisive final actions, for fine-tuning, without commitment to a final outcome. Grush’s proposal is that a mechanism for signal processing, the Kalman filter, performs this function. We can say that in using the Kalman filter the motor system is representing its goal.

Note that it is not obvious what would constitute “a representation” in this proposal. The Kalman filter is not a representation but a mechanism that allows representational activity by providing, among other things, appropriate isomorphism with the goal state. We could say that the Kalman filter, together with the motor-control system that employs it, constitutes a representation of the action in the static sense, but that would be inaccurate because the combined system as a whole is related to the goal action only when employed in the *process* of preparing for the final action. If we want to call that system a representation, then representations are much more ephemeral things than the language paradigm would suggest, and the philosophical interest in representations as

entities would lose its motivation. It is much more useful to take the motor-control system as a whole as a basic example of representational *activity*, because that will allow us to define “representation” in a firm, noncircular manner:

Representation is the process of performing goal-directed activity in a manner that allows the activity to be rehearsed and optimized in advance of the realization of the goal. This realization (whether planned or simply hypothesized) is what is represented by the activity.

Note that this definition includes the representational activity involved in perception, as Grush intends. He argues that perception involves anticipation of sensory input (sect. 5.3). With perception, the goal is the interpretation of sensory input, and the emulation system functions for that purpose as it does in the case of purely motor activity. This approach, holding that imagery constitutes anticipation of actual perception, has been proposed by Ellis (1995). Here again, Grush’s detailed account of the mechanism of the Kalman filter provides both support and testability.

The discussion of mental imagery is important because mental images have borne much of the weight of the representation-as-entity approach. Grush shows how mental imagery can be seen as a truncated or constrained version of representational activity, in which (using the Kalman filter) environmental sensory input is discounted and the representing process is “an off-line operation of an emulator” (sects. 4.4 and 4.5). Sensory input is replaced by internally generated input, and the initiating motivation is not actual perception but what Ryle (1949) would call “pretending” to perceive objects, knowing that one is not really perceiving them.

Conclusion. Of course the brain mechanisms for representation do not constitute art; they are natural and probably unconscious. But it is useful to see that artistic activity – creating representational objects for the purpose of inducing desired states in observers – can be understood as constituting the same functional activity as the motor paradigm. The artist, instead of presenting the viewer with a vase of flowers, allows the viewer to “rehearse” the activity of looking at a vase of flowers with a particular focus and affective response. Thus, the artist and the viewer are both engaged in representational activity. To say that the painting *itself* represents a vase of flowers can be a shorthand way of referring to the process, but unless this process is understood, in the way Grush has offered or in some similar way, as representation in the active sense, we will continue to wonder how one physical object can bear such a profound relationship to another physical object, with no visible connections and all by itself.

ACKNOWLEDGMENT

I am indebted to Ralph Ellis for stimulating discussions about the concept of representation.

Emulation theory offers conceptual gains but needs filters

Catherine L. Reed,^a Jefferson D. Grubb,^a and
Piotr Winkielman^b

^aDepartment of Psychology, University of Denver, Denver, CO 80208;

^bDepartment of Psychology, University of California at San Diego, La Jolla,

CA 92093-0109. creed@du.edu jgrubb@nova.psy.du.edu

pwinkiel@ucsd.edu <http://www.du.edu/~creed>

<http://www.du.edu/~jgrubb> <http://psy.ucsd.edu/~pwinkiel>

Abstract: Much can be gained by specifying the operation of the emulation process. A brief review of studies from diverse domains, including complex motor-skill representation, emotion perception, and face memory, highlights that emulation theory offers precise explanations of results and novel predictions. However, the neural instantiation of the emulation process requires development to move the theory from armchair to laboratory.

There is currently much interest, within psychology and cognitive science, to develop theoretical frameworks that integrate percep-

tual, motor, and conceptual processes in a common theoretical framework. Along with influential proposals by Damasio (1989), Barsalou (1999), and others, Grush is making a strong case for the importance of on-line and off-line simulation. His emulation theory extends previous work by emphasizing the role of efference copies and on-line, dynamic use of the feedback information. Further, he precisely specifies the characteristics of the control process and proposes some possible neural mechanisms. In our commentary, we focus mostly on the unique conceptual gains offered by the proposal, and highlight its fit to empirical data. However, we also suggest that more work is needed for the theory to achieve a respectable level of neurological plausibility.

Grush builds his main arguments around the motor system. Although the traditional simulation theories all highlight the general correspondences between the mechanisms underlying motor imagery and motor execution, the emulation theory, with its emphasis on on-line, efferent feedback, offers a more precise account of empirical data. This can be illustrated with a study that investigated the temporal relation between the physical and visualized performances of springboard dives (Reed 2002a). The study included participants across three skill levels (novice, intermediate, expert) and measured performance of different components within a dive (e.g., approach, hurdle). This design allows for a test of different predictions offered by the simulation and emulation accounts. The traditional simulation theory predicts that skill differences should manifest themselves only during the *first component* of the dive, during motor program selection. In contrast, emulation theory predicts that the skill differences continue throughout *all stages* of the dive, because the emulator would assess the consequences of each motor program selection as dives progress. Specifically, experts should use emulator corrections of their motor execution less because their motor programs are highly accurate and their selection is largely automatic. Novices cannot use the emulator for fine corrections because they simply lack motor programs with which to correct the errors. However, intermediates use the emulator the most to correct their selections of motor programs. The empirical data are consistent with emulation theory. Intermediate performers not only were comparatively slowest in the visualizations, but also showed predicted skill differences throughout the dives. Thus, emulation theory provides insight into the mechanisms underlying complex motor-skill imagery.

The gains offered by the emulation theory extend beyond the motor system. This can be illustrated by research on “embodiment” of emotion processing. Grush offers a useful idea that the “visceral/emotional emulator” helps not only in off-line processing (e.g., providing efferent feedback based on past decisions), but also in on-line processing of emotional material. Several findings not only support this general notion, but also highlight that the emulation process can be impaired on the “brain” level as well as the muscular level. Adolphs et al. (2000) observed that damage to right somatosensory-related cortices impairs recognition of emotion from facial expressions. Niedenthal et al. (2001) showed that participants required to hold a pen in their mouth (blocking efferent feedback) performed worse at detecting changes in emotional facial expression than participants allowed to mimic the expressions freely. Finally, a provocative study discussed by Zajonc and Markus (1984) found that participants who watched novel faces while chewing gum (motor-blocking condition) later performed worse on a recognition test than either participants who encoded by mimicking faces (muscular-facilitation condition) or participants who squeezed a sponge (motor-control condition). This finding is particularly important for the emulation theory because it shows that to benefit cognition, the emulator needs feedback from a specific effector, not just any sensory feedback.

There are many more findings like the ones just mentioned. In fact, a recent review of the social-psychological and emotion literature revealed a number of phenomena that can be explained by the ideas of simulation (emulation) and embodiment (Niedenthal et al., in press). In short, theories such as Grush’s, as well as Barsalou’s and Damasio’s, offer much promise not only when it comes

to accounting for specific problems in motor control, imagery, or emotion, but also as general theories of cognition, including social and emotional cognition.

However, to take advantage of the potential explanatory power of emulation theory, research must clarify how emulators are neurally instantiated. Grush ventures that the emulator for his prototypical system, the musculoskeletal system (MSS), is contained in the cerebellum. Since Grush’s theory explains imagery to be a product of emulation, this proposal predicts that damage to the cerebellum would disrupt motor imagery. However, in contrast to lesions in other brain regions such as the striatum, frontal lobes, and parietal lobes, cerebellar lesions are not known to induce deficits in motor imagery. For example, patients with Parkinson’s disease show selective deficits on motor imagery tasks, but patients with cerebellar atrophy do not (Reed & O’Brien 1996). This suggests either that emulation is not necessary for imagery or that there is more to MSS emulation than the cerebellum. In general, we suggest that emulators may not be discrete structures, but instead, capitalize on multiple subsystems of the brain.

In sum, emulation theory moves us beyond the current simulation theories by providing more mechanistic explanations and specific predictions. In its emphasis on the critical role of efferent feedback in efficient processing, emulation also gives current theories a good functional reason for why perception, cognitive, and motor systems are so tightly intertwined. Despite these strengths, the neural instantiation of such an emulator must be developed further so that it can incorporate multiple cognitive and motor functions. Further, the emulator, rather than being its own module, should be conceived as a general brain mechanism that permits feedback to multiple existing neural systems that have more or less direct relationships to the motor system. In short, the emulation theory requires some additional development before it fully filters down from armchair to laboratory.

Emulation of kinesthesia during motor imagery

Norihiro Sadato^a and Eiichi Naito^b

^aDepartment of Cerebral Research, Section of Cerebral Integration, National Institute for Physiological Sciences, Myodaiji, Okazaki 444–8585, Japan;

^bFaculty of Human Studies, Kyoto University, Sakyo-ku, Kyoto 606–8501, Japan. sadato@nips.ac.jp

eiichi.naito@neuro.mbox.media.kyoto-u.ac.jp

Abstract: Illusory kinesthetic sensation was influenced by motor imagery of the wrist following tendon vibration. The imagery and the illusion conditions commonly activated the contralateral cingulate motor area, supplementary motor area, dorsal premotor cortex, and ipsilateral cerebellum. This supports the notion that motor imagery is a mental rehearsal of movement, during which expected kinesthetic sensation is emulated by recruiting multiple motor areas, commonly activated by pure kinesthesia.

It is uncertain whether motor imagery could generate expected kinesthetic sensation, although it has been considered a mental rehearsal of movement. It is empirically known that many people can experience vivid motor imagery, mostly involving a kinesthetic representation of actions (Feltz & Landers 1983; Jeannerod 1994; Mahoney & Avenier 1987). In movement control, the musculoskeletal system is subject to the measurement of proprioceptive and kinesthetic information generated by actual movement and relayed as feedback sensory signals. One of the important predictions of Grush’s “emulation theory” in motor imagery is that the emulator will output the sensory signal in “mock” proprioceptive format in response to motor control signals (efferent copy), resulting in kinesthetic sensation. This is in contrast to the “simulation theory” in motor imagery, in which only efferent copies are supposed to be generated. If the emulation theory is correct, one may identify the output sensory signals generated by the emula-

tor by detecting their interaction with pure kinesthetic sensation without movement.

It is known that pure kinesthesia without movement can be elicited by vibration of the tendon with a specific frequency (83 Hz; Craske 1977; Goodwin et al. 1972a; 1972b; Naito et al. 1999). Using this fact, Naito et al. (2002) showed that motor imagery affected pure kinesthetic sensation, generated by tendon vibration without overt movement. They found that the motor imagery of palmar flexion, or dorsiflexion of the right wrist, psychophysically influenced the sensation of illusory palmar flexion elicited by tendon vibration. Motor imagery of palmar flexion psychophysically enhanced the experienced illusory angles of palmar flexion, whereas dorsiflexion imagery reduced it in the absence of overt movement. This finding indicates that the emulator, driven by the mental imagery, outputs the “mock” sensory signals in a proprioceptive format, which interferes with the real (but artificially generated) proprioceptive sensory information from the musculoskeletal system.

Another prediction of Grush’s emulation theory is that the articulated emulator is a functional organization of components (articulants), whose interaction is comparable to that within the musculoskeletal system, and hence their neural representations are expected to be common. This point was also demonstrated by the study of Naito et al. (2002). Regional cerebral blood flow was measured with O-15 labeled water ($H_2^{15}O$) and positron emission tomography in ten subjects. The right tendon of the wrist extensor was vibrated at 83 Hz (ILLUSION) or at 12.5 Hz with no illusion (VIBRATION). Subjects kinesthetically imagined doing wrist movements of alternating palmar and dorsiflexion at the same speed with the experienced illusory movements (IMAGERY). A REST condition with eyes closed was included. The researchers identified common active fields between the contrasts of IMAGERY versus REST and ILLUSION versus VIBRATION. Motor imagery and the illusory sensation commonly activated the contralateral cingulate motor areas, supplementary motor area, dorsal premotor cortex, and ipsilateral cerebellum. The researchers concluded that kinesthetic sensation associated with imagined movement was generated during motor imagery by recruiting multiple motor areas, which were also activated by the kinesthetic sensation generated by tendon vibration. These commonly activated areas may constitute the articulants of the emulator driven by the efferent copy during motor imagery.

In conclusion, generation of kinesthetic sensation during motor imagery, and its neural representation common to kinesthesia without movement, can be interpreted as “emulated kinesthetic sensation” in the framework of the emulation theory by Grush.

Modality, quo vadis?

K. Sathian

Department of Neurology, Emory University School of Medicine, WMRB 6000, Atlanta, GA 30322. ksathia@emory.edu

Abstract: Grush’s emulation theory comprises both modality-specific and amodal emulators. I suggest that the amodal variety be replaced by multisensory emulators. The key distinction is that multisensory processing retains the characteristics of individual sensory modalities, in contrast to amodal processing. The latter term is better reserved for conceptual and linguistic systems, rather than perception or emulation.

Grush develops his emulation theory as a unified account of perception, imagery, and motor control, with the prospect of extension to diverse other neural functions. This theory is an advance over previous, less systematic formulations of simulation and imagery as being important in sensorimotor function. It makes the claim that particular neural elements work together in an emulation of perceptual or motor tasks, running in a special mode in which they are disconnected from external inputs/outputs. Vari-

ous emulations differing in their characteristics can hence be run, based on which the organism can select the best one to implement in interaction with the environment. The appeal of the theory stems from its unifying potential, and hence its success will be measured to a large extent by how well its binding of seemingly disparate streams of thought bears up over time. In this commentary, I focus on the relationship between the proposed sensorimotor emulator and sensory modality.

Grush argues for modality-specific as well as amodal emulators in the nervous system. Modality-specific emulators are relatively easy to understand, in terms of the operation of modality-specific sensory or motor systems. For instance, the findings reviewed in section 4.3 of the target article are compatible with a role for a motor emulator during visual imagery. However, the concept of a strictly amodal emulator, one that is entirely independent of any sensory “tags,” is less clear.

Let me make it absolutely clear that I am not arguing against amodal representations in the brain. Such representations must exist for abstract concepts that can be encoded linguistically, or “propositionally,” rather than directly in the workings of sensory systems. Indeed, as a vital part of human thought and communication, they are among the most important abilities that evolution has conferred on our species, compared to the other species that live or have lived on this planet. It is the characterization of abstract, amodal representations as imagery, and, by extension, as substrates of emulation strategies, that I am not comfortable with. Rather than “amodal” emulators, I suggest invoking “multisensory” emulators to provide the link between modality-specific systems and between these systems and abstract representations. I must emphasize that this is not a merely semantic distinction. By “multisensory,” I mean a system that receives inputs from more than one sensory modality. The existence of multisensory processes is well established, as is their neural implementation. The functions of multisensory processing include integration between the senses, cross-modal recruitment of sensory cortical regions, and coordinate transformation. Each of these functions has been studied in some detail.

Coordinate transformation in multisensory neurons of the posterior parietal cortex (PPC) has been intensively studied by Andersen and colleagues. This work indicates that multiple reference frames are represented in different regions of the PPC (Buneo et al. 2002; Cohen & Andersen 2002; Snyder et al. 1998). Reference frames may be allocentric, as in Brodmann’s area 7a; eye-centered, as in the lateral intraparietal area (LIP) and parietal reach region (PRR); body-centered, as in LIP; and both eye- and hand-centered, as in Brodmann’s area 5. Further, the eye-centered neuronal responses in LIP and PRR are gain-modulated by a variety of other factors such as eye, head, body, or hand position (Cohen & Andersen 2002). This effectively allows for a distributed representation of multiple reference frames simultaneously, and hence, for the coordinate transformations that are required for particular tasks, for example, between the retinocentric reference frame of visual stimuli or the head-centered reference frame of auditory stimuli and the body-centered reference frame of reaching arm movements, so that motor outputs may be appropriately directed. Multisensory emulators, then, could be engaged for specific coordinate transformations to allow planning of motor behavior as dictated by the organism’s current goals.

Another function of multisensory neurons is to integrate perceptual processes across the different senses. Such multisensory integration has been studied at the level of single neurons in the superior colliculus (Stein & Meredith 1993) and more recently in human cerebral cortex using functional neuroimaging. A case in point is the integration of auditory and visual information during perception of speech, which appears to depend importantly on cortex in the superior temporal sulcus (Calvert 2001). Moreover, Freides (1974) suggested three decades ago that, regardless of the modality of sensory input, the task at hand, especially if it is complicated, will recruit the sensory system that is most adept at the kind of processing required.

One means of such cross-modal recruitment is imagery. For instance, visual imagery may accompany tactile perception and could play a role in the engagement of visual cortical areas during tactile perception. Such recruitment of visual cortex has now been demonstrated in a variety of tactile tasks involving perception of patterns, forms, and motion, and appears to be quite task-specific, with areas that are specialized for particular visual tasks being recruited by their tactile counterparts (Sathian et al. 2004). An alternative interpretation of this type of cross-modal sensory cortical activation is that the regions involved are truly multisensory rather than unimodal. There is, in fact, increasing evidence that cortical regions traditionally considered to be unimodal are actually multisensory, receiving projections from other sensory systems in addition to their “classic” sources (e.g., Falchier et al. 2002; Schroeder & Foxe 2002). Multisensory emulators could clearly be employed to facilitate such cross-modal recruitment and synthesis.

My point is that, in all these examples of multisensory and cross-modal processing, specific modality tags appear to accompany the relevant sensory representations, which are associated with corresponding coordinate systems. This differs from Grush’s account, in which there is an amodal system, devoid of specific modality tags, that is used for perception and for internal emulation. I suggest that such amodal, propositional systems are conceptual and linguistic rather than being perceptual or the substrate for either imagery or sensorimotor emulation. It will be important for future empirical and theoretical research to attempt to distinguish clearly between multisensory and amodal neural systems.

Brains have emulators with brains: Emulation economized

Ricarda I. Schubotz and D. Yves von Cramon

Department of Neurology, Max Planck Institute for Human Cognitive and Brain Sciences, 04103 Leipzig, Germany. schubotz@cbs.mpg.de
cramon@cbs.mpg.de

Abstract: This commentary addresses the neural implementation of emulation, mostly using findings from functional Magnetic Resonance Imaging (fMRI). Furthermore, both empirical and theoretical suggestions are discussed that render two aspects of emulation theory redundant: independent modal emulators and extra measurement of amodal emulation. This modified emulation theory can conceptually integrate simulation theory and also get rid of some problematic philosophical implications.

Emulators with brains. The emulation account provides a formal way to apply the idea that the brain’s default mode is not passive waiting but active prediction, not only in motor control and imagery, but also in perception and perceptual imagery – an extension which fits perfectly with a long series of fMRI studies we performed on voluntary anticipatory processes. These studies made use of the serial prediction task, which requires participants to predict perceptual events on the basis of stimulus sequences. The lateral premotor cortex (PM), pre-supplementary motor area (pre-SMA), and corresponding parietal/temporal areas are engaged in active anticipation of sensory events. Note that this network is activated in absence of motor behavior, and that perceptual input is controlled by contrast computation.

Several functional characteristics of the considered areas render them candidate components of an emulator network. First, in the aforementioned studies each PM field’s response is restricted to specific stimulus features: PM fields for vocal movements are engaged in rhythm and pitch prediction, those for manual movements, in object prediction, and those for reaching and pursuit, in spatial prediction. A simplified synopsis of the results indicates that the anticipation of sensory events activates the PM fields of those effectors that habitually cause these sensory events (Schubotz & von Cramon 2001; 2002; Schubotz et al. 2003). This

“habitual pragmatic body map” (Schubotz & von Cramon 2003) in PM may precisely reflect Grush’s description of an “articulated” body/environment emulator. Second, our findings would also be in line with an emulation network that entails both amodal and modal representations.

Grush proposes motor regions to reflect the controller, and ventral and dorsal processing streams to be the core environmental emulator. We would rather suggest that multiple PM-parietal loops (including the ventral/dorsal stream) function as emulators, with each loop linking both heteromodal and unimodal representations (following the terminology in Benson 1994). One may even hold articulated emulation to be the default mode of PM-parietal loops which are exploited for perception, action, and imageries (see Fig. 1). Visual, auditory, or somatosensory imagery might be generated by efferent signals to and feedbacks from the corresponding unimodal association cortices.

We argue that such a modal emulation cannot be considered to be independent from amodal emulation. Rather, the same signal is concurrently sent to both unimodal and heteromodal association areas, even though current internal and external requirements may then determine which feedback becomes causally effective. Visual, auditory, hand, and foot imagery may introspectively feel different possibly because the controller exploits different premotor-parietal-subcortical loops. But all these networks, first, are made of both unimodal and heteromodal cortices which, second, communicate with ease. Possibly this in turn renders an extra measurement process redundant, as we also argue. On the other hand, “controller” functions (or perhaps better, competitive filter functions) may be realized more restrictedly within pre-SMA, in turn under the influence of anterior median frontal cortices, lateral prefrontal cortex, and extensive feedback projections.

Don’t introduce independent modal emulators – even if imagery sometimes feels purely visual . . . An introspectively compelling reason for suggesting independent modal emulation is that some kinds of modal imagery (e.g., a vase) feel purely visual and not at all motor. However, our fMRI findings reveal introspective reports to be unreliable (because introspection does not tell us that motor areas are engaged in non-motor anticipation). Likewise, we are introspectively blind to the empirical fact that perceiving an object includes perceiving what is potentially done with that object (see Gibson [1979/1986] for the notion of an object’s affordance, and, e.g., Fadiga et al. [2000] for premotor responses to mere object perception in the monkey). Conversely, it is conceptually inconsistent to assume amodal emulators to be independent of modal emulators, because in the emulation account, perception is sensation, given an interpretation in terms of amodal environment emulators, whereas sensation in turn is the on-line running of modal emulators. It therefore appears that amodal and modal emulation have to be conceptualized as reciprocally dependent¹.

. . . And don’t measure the emulators – even if imagery sometimes feels proprioceptive. An introspectively compelling reason for suggesting extra measurement is that motor imagery feels proprioceptive and not at all dynamic/kinematic. This also builds the core premise for splitting emulation from simulation: A motor plan is a dynamic/kinematic plan, whereas full-blown motor imagery is (mock) proprioceptive by nature and therefore must be previously transformed from the former by intermediate emulation and measurement.² However, exactly this premise would be rejected by accounts based on the ideomotor principle (e.g., theory of event coding; Hommel et al. 2001). These take motor acts to be planned in terms of desired action effects, that is, expected sensory events, and therefore plans and effects most likely share a common neural code. Comfortingly, emulation theory is not committed to the view that efferent signals are motor by nature. To be an efferent signal is nothing more than to be a delivered signal, no matter whether motor, sensory, sensorimotor, or amodal. Let us assume that the controller speaks “Brainish,” the lingua franca spoken by every subsystem in the brain, and that “measurement” is nothing but (and therefore should be termed) feedback from

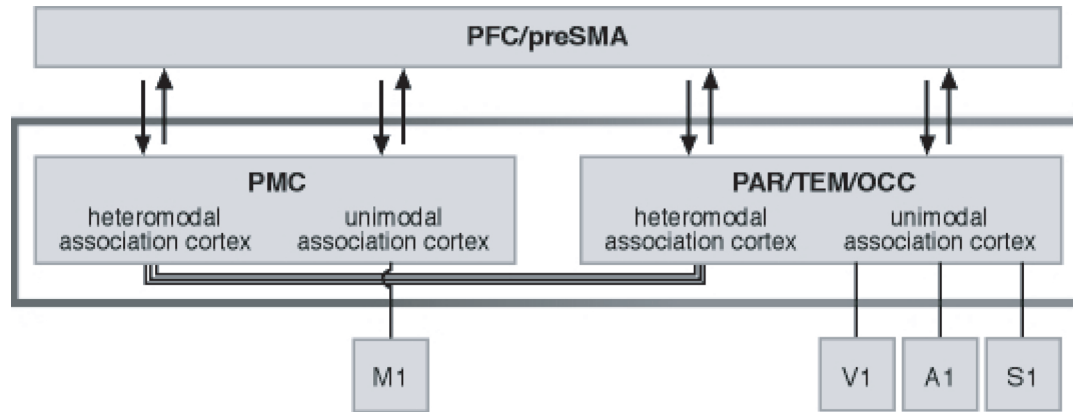


Figure 1 (Schubotz & von Cramon). Suggestion for a simplified neural implementation of a multipurpose emulator system (note that inverse models and/or Kalman filters might be added by cortico-ponto-cerebellar loops). Abbreviations: PFC, prefrontal cortex; PAR/TEM/OCC, parietal, temporal, occipital cortex; M1, primary motor cortex; V1/A1/S1, primary visual/auditory/somatosensory cortex, respectively.

the unimodal components of the general-purpose emulator. Grush correctly reminds us that “the emulator is a neural system: any and all of its relevant states can be directly tapped” (target article, sect. 5.1, para. 5).

ACKNOWLEDGMENTS

The preparation of this manuscript was supported by Marcel Brass, Christian Fiebach, Andrea Gast-Sandmann, Thomas Jacobsen, Shirley Rueschmeyer, Markus Ullsperger, and Uta Wolfensteller.

NOTES

1. Furthermore, redundancy emerges in an account that proposes one type of imagery (modal) to spring alternatively from two types of emulators (modal and amodal), but also two types of imagery (modal and amodal) to spring from one emulator (amodal).

2. Grush’s pilot-without-flight-simulator metaphor (sect. 3.1) exemplifies the necessity of measurement, but it also expresses how the measurement assumption suggests the introduction of little monolingual homunculi: A Turkish-speaking controller and a French-speaking emulator need a translation – the extra measurement.

Emulator as body schema

Virginia Slaughter

Early Cognitive Development Unit, School of Psychology, University of Queensland, Brisbane 4072, Australia. vps@psy.uq.edu.au
<http://www.psy.uq.edu.au/people/personal.html?id=35>

Abstract: Grush’s emulator model appears to be consistent with the idea of a body schema, that is, a detailed mental representation of the body, its structure, and movement in relation to the environment. If the emulator is equivalent to a body schema, then the next step will be to specify how the emulator accounts for neuropsychological and developmental phenomena that have long been hypothesized to involve the body schema.

Grush offers a detailed model of an information processing system that is hypothesized to represent the body as it moves in relation to the environment. His model of a body emulator appears to be consistent with the long-held notion of a “body schema.” The term “body schema” was introduced nearly 100 years ago in the neuropsychological literature as a unifying construct relating a number of disorders that seemed to indicate disturbances in the way patients perceived or conceived of their bodies (reviewed in Poeck & Orgass 1971). One syndrome that was originally cited as a disorder of the body schema was the phantom limb, which Grush explicitly mentions as explainable in terms of the operation of his hypothesized body emulator. This suggests that, at least

in some instances, the emulator and the body schema are equivalent.

The body schema was conceived as a dynamic representation of one’s own body, whose operation was most noticeable in cases of dysfunction (like phantom limbs) but whose normal functioning was thought to include a range of motor and cognitive phenomena from postural control to bodily self-concept. However, since the term was introduced, the notion of a body schema has become increasingly vague and overinclusive rather than more precise, and now there is genuine confusion over what a body schema is supposed to be and how it differs from a body image, body percept, body concept, or body awareness (for discussions, see Gallagher & Meltzoff 1996; Poeck & Orgass 1971; Reed 2002b).

Perhaps Grush’s emulator can bring some order to this confused state of affairs. He proposes that the emulator is a cognitive structure that represents the body as it moves and acts within the environment. This is at least part of what a body schema is supposed to be, and Grush’s model provides precise detail about how it may be instantiated. If the emulator and the body schema can be equated, this would represent a useful step forward because it can provide the field with a new and more precise definition of body schema, that is, a mental representation of the body that “implements the same (or very close) input-output function as the plant” (sect. 2.2), where the plant is the body being represented and controlled by the emulator mechanism.

Grush discusses how his emulator model can account for a number of low- and high-level cognitive phenomena, from motor control and motor imagery to simulation of another’s mental states. If we go ahead and equate the body schema with the emulator, then the model will also have to account for the neuropsychological phenomena that led investigators to postulate a body schema in the first place.

One of those is the phantom-limb phenomenon, which Grush addresses when he suggests that the existence of separate groups of phantom-limb patients who can and cannot move their phantoms can be explained by the amount of time the emulator experienced limb paralysis prior to amputation. This is a neat explanation. However, there are also reported cases of phantom limbs in congenital limb deletions (Poeck 1964; Weinstein & Sersen 1961) which are admittedly very rare but may pose a problem for the emulator model. Another neuropsychological disorder relevant to bodily motion and perception that the emulator should address is hemi-neglect, whereby the patient ignores perceptual information from one side of his visual field and as a result may stop using and lose a sense of ownership for his body parts on the affected side. Ideomotor apraxia is another classic disorder of the body

schema, defined as the inability to imitate gestures, which the emulator as body schema would have to explain.

Owing to imprecision in the way the term “body schema” was introduced, other disorders that have very different characters have traditionally also been classified as disturbances of the body schema. These include autotopagnosia (also called somatotopagnosia), whereby the patient becomes unable to recognize, localize, or name specific body parts, finger agnosia, whereby the patient is unable to recognize or name his fingers, and right-left disorientation, whereby the patient has difficulty identifying body parts as being located on the right or left side of their bodies.

Perhaps equating the body schema with the emulator will clarify the extent to which so-called disturbances of the body schema should be classified together. As several authors have already noted, phantom limbs and autotopagnosia are similar only insofar as they have something to do with the body; one is likely to be a low-level perceptual-motor dysfunction, whereas the other is a form of aphasia. If the emulator is equivalent to a body schema, then those disorders that do not implicate the specific cognitive machinery hypothesized to run the emulator can be reclassified, thereby bringing some order to an historically difficult-to-classify set of phenomena.

It should be noted that disorders of the body schema typically result from brain damage to the cortex, often the parietal lobes (Reed 2002b), but Grush hypothesizes that the emulator is located in the cerebellum (sect. 3.4).

The notion of a body schema most often appears in the literature on neuropsychological disorders, but it has also been proposed as the neurocognitive basis for neonatal imitation (Gallagher & Meltzoff 1996), whereby newborn infants match the bodily movements of adult modellers (usually facial movements but imitation of finger movements has also been reported; Meltzoff & Moore 1977). If the emulator is a body schema, can it account for neonatal imitation?

The emulator model that Grush presents is attractive because it is so precise, in stark contrast to the notoriously vague, ill-defined notion of a body schema. Although the emulator as presented cannot yet account for phenomena thought to implicate the body schema, there appears to be potential for the emulator model to bring some order to the neurological and developmental phenomena that have traditionally been thought to involve a body schema.

Evidence for the online operation of imagery: Visual imagery modulates motor production in drawing

Alastair D. Smith and Iain D. Gilchrist

Department of Experimental Psychology, University of Bristol, Bristol BS8 1TN, United Kingdom. Alastair.Smith@bristol.ac.uk
I.D.Gilchrist@bristol.ac.uk

Abstract: One property of the emulator framework presented by Grush is that imagery operates off-line. Contrary to this viewpoint, we present evidence showing that mental rotation of a simple figure modulates low-level features of drawing articulation. This effect is dependent upon the type of rotation, suggesting a more integrative online role for imagery than proposed by the target article.

Grush provides evidence for how imagery allows the individual to model the outcome of a particular action. This would necessarily operate off-line so that imagined movements are not carried out. Yet, if mental imagery is performed off-line, then a movement that results from this imagery system should be executed in the same manner as a movement made without being preceded by imagery. Preliminary data from our laboratory counter this viewpoint and suggest that mental imagery can actively (and unconsciously) modulate movement in a manner that is not predicted by the emulation theory of representation.

Models of drawing (Guérin et al. 1999; van Galen 1980; van Sommers 1984; 1989) typically assume the distinction between cognitive aspects of drawing and simple motor-output procedures. For example, the van Sommers (1989) model identifies a specified stage in the drawing process which controls the motoric components of the drawing procedure as a function of the constraints of both the movements involved and the materials used. For example, right-handers show a particularly strong tendency to begin their drawings at the top left of the figure. This bias is a highly robust feature of drawing production (Goodnow & Levine 1973; van Sommers 1984; Vinter 1994). Along with other graphic production rules (e.g., pen lifts), starting position is determined to optimise drawing efficiency in light of movement restrictions (Thomassen & Tibosch 1991). As such, they should not be affected by the nature of the more central cognitive processes that precede them.

To test the extent to which imagery occurs off-line, and prior to graphic production rules, we compared simple figure copying with production of the same figure following a mental rotation. The graphic output is identical in both conditions, but the preceding cognitive processes differ. As a result, any difference in drawing articulation between conditions must be due to an online influence from imagery. Participants copied right-angled triangles, which are known to have a highly stereotypical starting position (van Sommers 1984). On each trial we recorded the vertex at which participants began drawing.

In Experiment 1, participants ($N = 24$) were presented with sheets of A5 (210 x 148 mm) paper on which was printed a single right-angled triangle in the centre of the upper portion of the page. Stimuli measured 2 cm in height and width, and could be at one of eight orientations (after van Sommers 1984). Participants were required to mentally rotate the figure either 90° clockwise or 90° anticlockwise. They then drew the product of their mental manipulation below the original figure. Starting positions were compared with those in a copy condition for the same triangles. Participants began at a vertex that was consistent with a *motor* response (i.e., starting at the same vertex as when producing the same figure in the copying condition), an *imagery* response (i.e., starting at the same vertex as when copying a pre-rotated version), or *neither* type of response.

Start-point frequency data were transformed into log-odds ratios (see Wickens 1993) and analysed using one-sample *t*-tests. There was a significantly greater number of *motor* responses than *imagery* responses ($t = 13.9$, $df = 23$; $p < .0001$). More crucially however, there was a significantly greater number of *imagery* responses than *neither* responses ($t = 4.04$, $df = 23$; $p < .001$). This shows a reliable effect of mental rotation on starting position: starting position was rotated with the triangle. In addition, participants showed the same number of pen lifts for the copy and rotation conditions, and comparable error rates.

Starting position was not completely subject to low-level peripheral factors (subsequent to imagery). Instead it seems to be determined prior to, or during, the manipulation of the image. This argues against a strictly serial account of drawing production (van Sommers 1989) and also suggests that mental imagery is not as off-line as the emulation theory proposes. As Grush points out in the target article, there is increasing evidence to suggest that image manipulation is intrinsically linked to the motor system, even when simply planning a rotational movement (Wohlschläger 2001). It also seems to be the case that the way the mental rotation is conceptualised affects the likelihood of motor structures being used to support the manipulation (Kosslyn et al. 2001). Wraga and colleagues (2002) found that when participants imagined rotating a stimulus there was activation of primary motor cortex, but when they imagined themselves rotating around a stationary stimulus there was no motor activation.

Following on from this, in Experiment 2, participants were asked to imagine *themselves* moving around the page to view the triangle from the point denoted by an arrow marker (i.e., with the arrow pointing to what is now the bottom of the stimulus). When

there was no arrow present, participants reproduced the stimulus without mental manipulation. This manipulation resulted in the same drawing production as in Experiment 1, but through a different conceptualisation of the rotation.

In this experiment there was no evidence for the rotation of starting point consistent with the metal rotation. There was a significantly greater number of *motor* responses than there were *imagery* responses ($t = 13.7$, $df = 23$; $p < .0001$). However, there was a significantly greater number of *neither* responses than there were *imagery* responses ($t = -2.44$, $df = 23$; $p < .05$). The results from Experiments 1 and 2 together suggest that the manner in which participants transform mental images can modulate the likelihood of an influence on motor articulation. In turn, these results suggest that some mental imagery processes may be more off-line than others (cf. Kosslyn et al. 2001).

It seems that mental imagery is not an entirely off-line process. Depending on the nature of the transformation, image manipulation can have low-level effects upon movement articulation. Participants are largely unaware of such effects (see Vinter & Peruchet 1999), and this contrasts with the highly goal-driven form of visuomotor imagery that Grush discusses. In his article, Grush tacitly takes the position that mental imagery is a cognitive function: although there may be concomitant activation of relevant motor or perceptual structures, it mainly serves to drive imagery through the provision of efference copy. Our data point towards a more integrated and ubiquitous role for mental imagery, which does not operate in isolation but in a more dynamic and interactive manner.

If emulation is representation, does detail matter?

Lynn Andrea Stein

Computers and Cognition Laboratory, Franklin W. Olin College of Engineering, Needham, MA 02492-1245. las@olin.edu
<http://faculty.olin.edu/~las>

Abstract: Grush describes a variety of different systems that illustrate his vision of representation through emulation. These individual data points are not necessarily sufficient to determine what level of detail is required for a representation to count as emulation. By examining one of his examples closely, this commentary suggests that *salience* of the information supplied is a critical dimension.

Flattered as I am by Grush's reference to my work, I fear that he has conflated certain of my efforts with the research of a colleague with whose original system I began. I think that it is worth taking a closer look at our projects, not merely to set the record straight but also because it will shed some light on the nature of emulation as a representation system.

Toto was a robot, built by Maja Mataric, that was capable of randomly wandering the corridors and lounges of the MIT AI Lab (Mataric 1992). In addition, as Toto wandered through this space, it recorded – in a clever and innovative pseudo-representational way that was largely the point of Mataric's project – the gross features of space that it had encountered: Wall Left, Corridor, Open Space, and so on. Together with annotations regarding transitions, this “memory” of where Toto had previously been allowed it to return intentionally to a particular space. If I understand Grush correctly, Toto's “memories” served the robot as a kind of abstract, perhaps unarticulated, emulator of Toto's navigational behavior against which its future goal-directed behavior could be measured and – in a closed loop – driven.

What is particularly nice about regarding these memories as emulation of the world is the contrasting level of detail at which actual sensory and motor data exist versus the gross generalizations of memory representations like Corridor. That is, it is not at all necessary for an emulation to preserve all of the detail of the

actual operation of the robot plant; it needs merely to track the *salient* aspects of that operation, in this case whatever data are sufficient for place recognition and prompting of where to turn. This requirement of salience rather than precision frees emulation to operate as a sort of abstractor, folding together all of the possible ways to roll down the central hallway into the single abstract memory, Corridor.

Mataric's work on Toto provides one set of insights into emulation as representation. My research went in another direction entirely. Like Grush, I was interested to know how far this kind of implicit representation could scale. I observed that Toto could navigate to specific places, but only after its emulator-memory had been trained up by prior experience with that location. My question was whether Toto could be made to go to *new* places, places of which it had only been told. My solution was to use this novel information to feed the emulator, programming it up to “remember” places that Toto had never been (Stein 1994).

In order to accomplish this, I exploited a fundamental fact of Toto's architecture: The best trainer of the emulator/memory is experience, and so in my augmented system – called MetaToto – the robot learned by actually experiencing these hypothetical locations. This, in turn, involves something that I called imagination: (Meta)Toto in essence hallucinates wandering through a world that is described and builds ersatz memories of these places. Here, though, it is the actual robot brain's sensory and motor-control systems – excepting only the final layer that perceives or acts in the real world, which is temporarily disconnected – that do the actual work of training up the memory/emulation system.

Perhaps, however, Grush would not count this hallucinatory imagining as emulation. After all, it is not being used to correct or project the activity of the robot control system. Further, the information that is used to create this hallucinatory experience is not learned but provided directly from the description.

If true, this is an ironic turn of events. Although it is entirely external to the robot brain, the hallucinatory experience is exactly an emulation – in the classical, if not Grushian, sense – of the robot's actual would-be experience in the real world. And it is articulated in essentially the same way as the robot's actual sensorium and (to a lesser extent) motor apparatus. So perhaps I misunderstand Grush, and he would accept this hallucinatory experience as exactly the kind of emulation system he is proposing. Indeed, it bears more than a passing resemblance to his Figure 7.

Grush does not really say how closely his emulators need to track the actual musculoskeletal system (MSS), although he does use the idea of articulation to give some sense of where he thinks the major similarities must lie. In Toto, memory provides a very abstracted representation of past action, sufficient to guide future navigation but far from definite or determined in the way that a motor plan would be. In MetaToto, the hallucination of moving around in an imagined environment is also inaccurate, but the articulations of this emulation (if emulation it is) are much more like those of (Meta)Toto's own sensorium.

All of this raises the question of what, exactly, an emulator is. Clearly it is something that maps from actions (or action commands) to expected sensations, modeling the behavior of the (body and) the world. But Toto's emulation, in the form of memory, tracks only gross properties of space – Corridoriness, for example – whereas MetaToto's hallucination actually supplies (imaginary) readings for each of Toto's 12 sonars. So maybe the concrete/abstract dimension can vary, and what is really important is salience: providing the articulations that are necessary for whatever behavior the emulation will support.

Representation: Emulation and anticipation

Georgi Stojanov^a and Mark H. Bickhard^b

^aComputer Science Department, Electrical Engineering Faculty, Saints Cyril and Methodius University, 1000 Skopje, Macedonia; ^bCognitive Science, Lehigh University, Bethlehem, PA 18015. geos@etf.ukim.edu.mk
mark@bickhard.name <http://bickhard.ws/>

Abstract: We address the issue of the normativity of representation and how Grush might address it for emulations as constituting representations. We then proceed to several more detailed issues concerning the learning of emulations, a possible empirical counterexample to Grush's model, and the choice of Kalman filters as the form of model-based control.

We are quite sympathetic with much of the general orientation involved here – in particular, the emphasis on anticipation – and we would like to explore a few important issues that are embedded in Grush's discussion.

Grush begins his discussion with an engineering perspective on motor control, but later makes use of such notions as “expectation” and “representation.” This shift crosses a major divide, that between fact and norm. The normativity in this case is that of the truth value of representation, of the capability of being true or false. In general, when one discusses phenomena such as representation, there are issues that cannot be avoided, such as normativity and learning. In our commentary we will focus on these broader issues.

Is there any concern here for “emulations as representational” being able to model the possibility of being true or false? If there is no such concern, then in what sense is representation involved at all? The shift into this language engages these issues and makes use of the special normative properties of representation in the discussions of more cognitive phenomena. Without any representational normativity, much of the last part of the target article makes no sense, so there would seem to be an embedded necessity for ultimately addressing these issues.

If there is a concern regarding representational truth value, then is the issue one of being true or false from the perspective of the organism, or just in terms of the usefulness of an external observer making representational attributions to the organism? If there is such a concern and it is from the perspective of an observer (e.g., Clark 1997; Dretske 1988), then how does one account for the representations of the observer – isn't the observer just one more homunculus, even if an external one? An attempt to render representation as strictly constituted in the ascriptions of external observers seems to derive all of its representational normativity from those observers themselves, and, therefore, does not accomplish any sort of naturalistic model of that normativity.

If representational truth value is to be modeled from the perspective of the organism, then how does one make good on the normativities involved in “true” or “false”? How does one cross Hume's divide? One way to begin to address such normativities would be in terms of the anticipations involved in the emulations: Such anticipations can be true or false, and can, in principle, be detected to be true or false by the organism itself. In such a model, the normativity of representation is derived from the functional normativity of anticipation, which then must itself be accounted for. A model of representation and cognition based on such functional anticipations has been under development for some decades (e.g., Bickhard 1980; 1993; 2004; Bickhard & Terveen 1995). So, a question for Grush: Do you have an approach for addressing these issues, and, if so, could you outline it a bit?

A cautionary note: One powerful way to attempt to model the normativities of function, including, potentially, the function of anticipation, is in terms of the etiology (generally the evolutionary etiology) of the part of the organism that has the function at issue (Millikan 1984; 1993). The kidney, in such an approach, has the function of filtering blood because the evolutionary ancestors of

the kidney were selected for having that effect. Unfortunately, this does not work as a naturalized model of function. Here is one reason: Millikan points out that, if a lion were to magically pop into existence via the coming together of molecules from the air, that lion is, by assumption, molecule by molecule identical to a lion in the zoo; nevertheless, the organs of the lion in the zoo would have functions, whereas those of the science fiction lion would not because it does not have the right evolutionary history (Millikan 1984; 1993). What this example points out, however, is that etiological function cannot be defined in terms of the current state of the organism, but only current state, according to our best physics, can have proximate causal power. Here we have two lions with identical dynamic properties, but only one has functions. Etiological function, therefore, cannot make any dynamic difference – any causal difference – in the world, and therefore, does not constitute a successful naturalization of function (Bickhard 1993; 2004; Christensen & Bickhard 2002). In consequence, then, the etiological approach to the normativities of representation is similarly blocked.

We turn now from issues of normativity to some more detailed and empirical considerations. First, we note that Grush does not discuss the issue of learning “emulations as representations.” If we pursue the control theory framework, we might, for example, look for system-identification tools for obtaining the relevant Kalman filter (KF) parameters. What are their biological counterparts? Further, the computations in KF can be seen as manipulations of multivariate normal probability distributions – what guarantees that those conditions are fulfilled? Can it be shown that the approximations involved are good enough?

In the context of situated robotics, Grush mentions the Meta-Toto architecture of Stein as being able to engage in off- and on-line use of the map it builds of the environment in order to solve navigation problems. MetaToto is built on the basis of Mataric's Toto architecture (Mataric 1992). Is there any essential difference with respect to emulation versus simulation between these two architectures?

When it comes to empirical (counter) evidence, we would like to mention O'Regan's change-blindness experiments (e.g., O'Regan & Noë 2001). The setup is quite simple: A subject is shown a photo in which suddenly some mud splashes (or flickering) appear, and in the meantime a drastic change (up to one third of the size of the overall picture) is introduced. The majority of the subjects do not notice the change. According to Grush's KF framework, because the estimate does not match the stimuli, the Kalman gain should increase, which would lead to an accurate representation and perception of the changed photo. Apparently some crucial elements dealing with attention are missing from the proposed framework.

Finally, we have a question concerning the role of emulation in general, and KFs in particular, in the overall model. Concerning KFs, it would seem that any model-based control theory could do what Grush needs here – is there a more specific reason to choose KFs? In this respect, see some of our works and papers cited here (Stojanov 1997; Stojanov et al. 1995; 1996; 1997a; 1997b). Concerning emulation more broadly, the initial motivation for the model is in terms of motor-control emulators. These depend on, among other things, explicit efferent and afferent transmissions with respect to an emulator process. But the later uses made of the notion seem to depend on more general notions of modeling and anticipations generated from such modeling. The question, then, is whether other forms of generating anticipations – such as, for example, the set-up (microgenesis) of an interactive system to be prepared to handle some classes of interactions but not others, thus anticipating that the interaction will proceed within that anticipated realm (Bickhard 2000; Bickhard & Campbell 1996) – might not do as well for broader forms of cognition?

In sum, we are quite enthusiastic about the modeling orientation that Grush has discussed, but we would contend that some of the important issues have yet to be addressed – at least in this paper.

A neuropsychological approach to motor control and imagery

Barbara Tomasino, Corrado Corradi-Dell'Acqua,
Alessia Tessari, Caterina Spiezio, and Raffaella Ida Rumiati
*Cognitive Neuroscience Sector, SISSA/ISAS International School of
Advanced Studies, Trieste, 34014, Italy. tomasino@sisssa.it
corradi@sisssa.it tessari@sisssa.it spiezio@sisssa.it
rumiati@sisssa.it <http://www.sisssa.it/cns/npsy/neuropsych.htm>*

Abstract: In his article Grush proposes a potentially useful framework for explaining motor control, imagery, and perception. In our commentary we will address two issues that the model does not seem to deal with appropriately: one concerns motor control, and the other, the visual and motor imagery domains. We will consider these two aspects in turn.

The author takes into account many possible types of feedback that can be used to guide movement planning. However, the important effect of visual hand/limb position during motor tasks (visual feedback) is not taken into consideration by the author. But many studies on both healthy individuals and patients have shown a consistent effect of visual feedback on task performance when participants could see their own hands/limbs (Botwinick & Cohen 1998; di Pellegrino et al. 1997; Graziano 1999; Maravita et al. 2003; Pavani et al. 2000). Moreover, a similar effect has been found also to be at work in motor imagery tasks (Sirigu et al. 2001).

The proprioceptive/kinaesthetic feedback is another example discussed by Grush. The author suggests that the phantom-limb phenomenon is a clear sign of the presence of an emulator, which corrects its input-output configuration after the limb loss. Grush argues that phantom-limb patients, who suffered a long period of pre-amputation paralysis, have time to recalibrate the emulator, leading to a deficit in imaging movements of the paralyzed limb. This argument however is not consistent with the Brugger et al. (2000) report of a patient with a congenital hand deletion who showed an intact ability to imagine movements of both hands even though he had never experienced a feedback. Likewise, in discussing the distinction between simulation and emulation, Grush completely neglects the case of ideomotor apraxic patients (de Renzi et al. 1980; Liepmann 1905), who "know" the motor plan of actions they cannot implement into proprioceptive/kinaesthetic outputs.

The model is also applied by Grush to the visual and motor imagery domains. To demonstrate that the model can account for phenomena in the visual imagery domain, Grush mentions tasks that we argue are not only visual. In particular the author reports, as an example of a visual imagery operation, the ability to predict the consequences of an action in the visual scene. However, this mental operation (i.e., mental rotation) is not simply visual but requires a motor transformation (Kosslyn et al. 2001). Hence, mental rotation relies on a network that is distinct from that used in generating and inspecting images (Cohen et al. 1996; Kosslyn et al. 1993; Parsons & Fox 1998). The tasks that are known to tap visual imagery without involving a motor component, such as the Island Test (Kosslyn et al. 1978), the Clock Test (Grossi 1993), the Piazza del Duomo (Bisiach & Luzzatti 1978), have not been considered by the author. Is his omission due to the fact that the classical approach to visual imagery would not easily fit the model?

Furthermore, Grush mentions two main findings that according to him should prove the validity of the emulation model applied to the visual imagery domain: (1) a general overlapping activity between the visual areas activated during overt perception and visual imagery tasks; and (2) a huge similarity (i.e., isomorphism) of brain activity in both overt and imagined perception (sect. 3.2). However, overlapping brain networks and isomorphism are properties that, rather than being distinguished features of the emulation theory, are common to both the simulation and the emulation theories. Earlier in the target article (sect. 2.3), the author states that the simulation theory itself is not sufficient to explain the motor imagery phenomena and claims that an emulator of the musculoskeletal system is needed. When he then turns

to discuss the visual imagery domain, it becomes far from clear where the simulation ends and the emulation starts, raising the doubt whether the model is applicable outside the motor domain.

Our final comment refers to the neural bases of motor imagery. Among those areas activated during both execution and imagination, Grush mentions the importance of premotor and supplementary motor areas as well as the cerebellum. However, there are several neuroimaging (Decety et al. 1994; Grafton et al. 1996; Stephan et al. 1995) and neuropsychological studies (Rumiati et al. 2001; Sirigu et al. 1996) indicating that the superior or inferior parietal cortices are also critically associated with imagination and execution of a movement. In addition, we disagree with Grush's assertion that only the primary motor cortex (M1) "is conspicuously silent during motor imagery" (sect. 3.2, para. 2) among the brain regions supporting the simulation of movements. In fact, whereas there are many other studies providing reliable evidence that M1 is actually involved in the simulation of a motor act (e.g., Decety et al. 1994; Gerardin et al. 2000; Kosslyn et al. 2001; Lang et al. 1996; Porro et al. 1996; Roland et al. 1980; Stephan et al. 1995), there is only one study consistent with Grush's claim (Richter et al. 2000). M1 not only is an efferent motor area, but is also involved in processing higher cognitive functions, as has been shown by neurophysiological studies in monkey (Alexander & Crutcher 1990; Ashe et al. 1993; Carpenter et al. 1999; Georgopoulos et al. 1989; Pellizzer et al. 1995; Smyrnis et al. 1992; Wise et al. 1998), functional neuroimaging studies (Catalan et al. 1998; Grafton et al. 1995; 1998; Karni et al. 1995; 1998; Lotze et al. 1999; Tagaris et al. 1998), and experiments using the transcranial magnetic stimulation technique in humans (Chen et al. 1997; Ganis et al. 2000; Gerloff et al. 1998; Tomasino et al., in press). More precisely, it has been shown that M1 plays a role in stimulus-response incompatibility, plasticity, motor sequence learning and memory, learning sensory-motor associations, motor imagery, and spatial transformations.

In conclusion, we appreciate the intellectual effort made by the author in proposing a model that deals with a wide range of cognitive phenomena. However, the modelling would be more effective had the author identified more carefully the computations involved in the tasks tapping the various domains he considers, and had he related them to the neural mechanisms more thoroughly.

Sensation and emulation of coordinated actions

Charles B. Walter

*School of Kinesiology, University of Illinois at Chicago, Chicago, IL 60608.
cwalter@uic.edu*

Abstract: Although the application of the emulation model to the control of simple positioning movements is relatively straightforward, extending the scheme to actions requiring multisegmental, interlimb coordination complicates matters a bit. Special consideration of the demands in this case, both on sensory processing and on the process model (two key elements of the Kalman filter), are discussed.

Feed-forward, feedback, and efference copy mechanisms have long served as cornerstones of motor-control models. In a very accessible account, Grush provides a formal synthesis of these three mechanisms that exploits the advantages of each. Following Wolpert et al. (1995), a Kalman filter is incorporated into the control scheme. Although I found the discussion and speculation concerning potential applications of this model to imagery, perception, and other cognitive functions fascinating, I will limit my comments to issues of motor control.

It is first perhaps relevant to note that some have criticized the development of human motor-control models based on multiple mechanisms as playing a rich man's game of science. The alternative view adopted here is that the clear advantages of such mod-

els and evidence of redundant control inherent in many biological systems justify this approach. The next step from this perspective is to elaborate each process in the model in disparate contexts, a task receiving somewhat preliminary treatment in the article because of space limitations. Two of the processes are discussed here with respect to movement coordination.

Perhaps unsurprisingly, simple feed-forward and feedback models of motor control have often (though certainly not exclusively) been tested using single-limb, point-to-point movements. The control of even these relative “simple” actions is far from trivial, but the function of each component in Grush’s control scheme is relatively straightforward in this case. But what of more complex, multi-limb actions that implicitly entail coordination per se as part of the goal? Examples range from maintaining a 180-degree interlimb phase offset when beating a drum or when walking, to generating the constantly changing (but determinate) phase of the notes constituting a jazz riff on piano. The model still potentially applies in this case, but the processing demands of some model components are naturally increased. This is particularly the case for the measurement matrix in the “plant” (i.e., the record of appropriate sensory state variables) and the process model in the emulator.

Consider, first, the appropriate sensory variables constituting the measurement matrix for the state of the plant during coordinated actions. Different classes of sense organs provide, to lesser and greater degrees of accuracy, unidimensional information regarding muscle length and change in length, joint angle, muscle tension, and so forth. A direct measure of phase offset is obviously not available. Discrete point estimates of relative phase for drumming at a constant rate could be determined by monitoring the consistency of the half-periods between bilateral contact points, for example. But a *continuous* measure of phase lag, which is required to exploit the rapid feedback afforded by the model, requires substantially more processing; determining the phase offset between the end points of multisegmented effectors requires still more processing; and so forth. The problems are not intractable, but transforming the raw information into a form relevant to the task at hand presumably requires time and is subject to error.

Now consider the complexity of the process model that is needed to emulate a coordinated action. A number of mappings reflecting the mechanics associated with single-limb control must be included, comprising such considerations as the length-tension and force-velocity properties of muscle, gravity, passive torques interacting nonlinearly among segments, and so forth. When additional limbs are involved in the action, the need to consider so-called “coordination dynamics” is invoked. As Kelso and colleagues noted some time ago (e.g., Haken et al. 1985), concurrent limb movements are influenced not only by a central control signal but often by nonlinear coupling between the effectors as well. This coupling can result in (initially) unpredictable behavior. A transition to a completely unintended phase offset between the limbs can occur if the drive is scaled past a critical level, for example. The timing demands of the effectors clearly influence interlimb interactions, but spatial interference is evident as well (e.g., Walter et al. 2001). The dynamics of these space-time interactions are part of more global plant dynamics, so the model is not invalidated. The obvious point is simply that the process model for human coordination is substantially more complex than, for example, the mapping between the signal driving torque motors and the resultant motion of a robot arm. But the more insidious question that emerges when attempting to fill in the details of this model component is whether the complexity of the emulation process is amenable to our attempts to capture it with an elegant, articulated account.

Although presented separately, the twin concerns elaborated above are of course related. The Kalman gain is only as effective as the rapidity and accuracy with which the actual sensory signal and the mock sensory signal are generated. Illuminating these two processes alone poses an interesting challenge in further developing the model to account for complex, coordinated actions.

Representing is more than emulating

Hongbin Wang^a and Yingrui Yang^b

^aSchool of Health Information Sciences, University of Texas Health Science Center at Houston, Houston, TX 77030; ^bDepartment of Cognitive Science, Rensselaer Polytechnic Institute, Troy, NY 12180.

Hongbin.Wang@uth.tmc.edu yangryi@rpi.edu
<http://www.shis.uth.tmc.edu/>

Abstract: Mental representations are more than emulations. Different types of representations, including external representations, various mental models (distorted and abstract), and emulative models, can all play important roles in human cognition. To explain cognitive performance in a specific task, a systematic analysis of the underlying representational structures and their interactions is needed.

The essence of the target article is that the human brain maintains various internalized and adaptive representations of the body and the environment, which, via Kalman filter–like control schemes, can emulate/simulate/mimic quite precisely the actual brain-body and human-environment interactions. Although the idea of mental representations as emulations is intuitively appealing and theoretically interesting, we would like to comment on the adequacy of the theory and argue that mental representation is more than emulation.

Few will doubt that a mental emulator would be useful if one exists. As demonstrated extensively in the target article, an emulator implements an essential look-ahead capacity and therefore permits a systematic treatment for a wide variety of cognitive phenomena, including motor and visual imagery, perception, reasoning, and language. However, the real issue is whether these powerful mental emulators do exist in the brain and in the mind. We would expect that an emulation theory of representation would provide answers to important questions such as why mental emulators are necessary and how they are acquired. Unfortunately, although the target article provided a detailed account of how mental emulators help to implement various cognitive functions, it has left these fundamental questions largely unanswered. The author has felt reluctant to give evidence that *counted against* the theory and admitted that the goal was not “to provide compelling data for its adoption” (sect. 6.4, para. 1). The result is a less convincing theory.

In this commentary we raise two issues that we think are directly related to the theory’s treatment of mental representations as emulations. The first one has to do with the relationship between emulation and other types of representations. Numerous studies in the broad field of cognitive psychology and cognitive science have shown that people can and do use different types of representations in performing various cognitive tasks (Donald 1994; Palmer 1978; Rumelhart & Norman 1988). To explain cognitive performance in a specific task, an analysis of the underlying representational structures and their interactions is needed. For example, Zhang and colleagues have systematically studied the nature and function of external representations (Zhang 1997; Zhang & Norman 1994). According to this view, certain types of information, such as affordance and salient spatial relationships, exist as external representations in the world and can be directly perceived and used by the mind. As a result, for these types of information no mental internalization is needed to support actions. In one of the studies, Zhang et al. found that the subjects’ performance in solving different isomorphs of tic-tac-toe was determined by the interplay of different types of representation and that, when more information was available directly through external representations, the task was easier.

Although various internal representations are an important type of representations, they are typically compressed, segmented, and distorted forms of the represented entity but not emulations. These characteristics result from the mind’s nontrivial reorganization of sensory input and are constrained by the limited processing power of the brain. One example of this is how the mind represents space (Tversky 2000). Recent evidence in cognitive neuroscience has shown that the “where” pathway is more complex than a unified object-location emulator. It consists of multi-

ple spatial centers, each processing and extracting distinct spatial information. As a result, space is represented in the mind not once but multiple times, not in a unified manner but segmented. Each representation is a (distorted) salience map with distinctive frames of reference (see Wang et al. 2001).

While we agree that emulation is yet another type of mental representation, we maintain that it is the goal but not the means. After all, when we have an emulator of the universe, to a certain degree we have understood how it works (see also Lewontin 2001 who suggests that the essence of science is to seek metaphor/emulation). In early development children are capable of interacting with the environment, but it is implausible to assume that they possess an emulator of the environment. Learning to interact with a model/emulation of the environment is certainly desired, but children must start with a model-less interaction mode and gradually learn to acquire such a model (e.g., Sutton & Barto 1998).

We would like to clarify our argument further by raising the second issue, which is related to the author's attempt to use amodal environment emulation to explain Johnson-Laird's mental models in human reasoning. First of all, according to Johnson-Laird (e.g., Johnson-Laird 2001), each mental model represents a possibility, and people by default represent only what is true but not what is false. In this sense each mental model is a fragmented representation of what is true, and constructing mental models involves considerable information processing at the semantic level. Treating mental models as emulations raises questions about how the emulators are smart enough to perform these semantic processes. Therefore, it seems that the claim does not solve the problem but simply gives it a different name.

In addition, treating mental models as space/object emulations against sentential representations would encounter difficulty in representing certain types of ordinary inferences that involve compound statements (Yang & Bringsjord 2003). The current mental model theory is powerful in representing prepositional reasoning but it has limited power in quantified predicate reasoning. Mental models theory provides initial models for truth connectives such as disjunction, and provides initial mental models for quantified atomic sentences such as: "(For all x) Ax." In combination, it can represent a compound statement such as: "(For all x) Ax or (for all x) Bx." It is difficult, however, for the current mental models theory to have a spatial/object style presentation for a quantified compound statement such as: "(For all x)(Ax or Bx)." For statements of this type, Yang and Bringsjord (2003) argue that some special syntactic structures would have to be constructed. Though one may still find ways to represent these mental models as space/object emulations (e.g., $[x]_{Ax, Bx}^{Ax}$, as a quantified disjunction), they apparently make mental models less efficient. At this point, the emulation hypothesis would face a philosophical question: Does cognition need to be maximally efficient?

Small brains and minimalist emulation: When is an internal model no longer a model?

Barbara Webb

Institute for Perception, Action and Behaviour, School of Informatics, University of Edinburgh, Edinburgh EH1 2QL, United Kingdom.
bwebb@inf.ed.ac.uk
http://www.informatics.ed.ac.uk/people/staff/Barbara_Webb.html

Abstract: Many of Grush's arguments should apply equally to animals with small brains, for which the capacity to internally model the body and environment must be limited. The dilemma may be solved by making only very approximate predictions, or only attempting to derive a "high-level" prediction from "high-level" output. At the extreme, in either case, the "emulation" step becomes trivial.

Grush makes a good case that concepts of emulation from control theory can illuminate a range of issues when applied to the brain.

But this approach should apply to a wider range of brains than the human ones that form the focus of the article. Animals with small brains, such as insects, face many of the same behavioural challenges that the emulation theory is proposed to solve. For example, flies need to distinguish the visual slip that results from self-initiated steering, from that caused by external disturbances if they are to use the latter for stabilisation. It was just this sort of problem that first led to the explicit formulation of the concept of efference copy by von Holst and Mittelstaedt (1950; simultaneously described by Sperry 1950 as "corollary discharge").

Grush notes, in the field of motor imagery, a surprising lack of appreciation of the fact that an internal copy of the motor output is not in itself sufficient to explain the prediction of sensory experience. This failure to deal with the problem of the difference between the format of the output and input signals is pervasive in the biological literature on efference copy. Von Holst and Mittelstaedt, in their early formulations, describe efference copy as being compared to the sensory reafference like a "photo-negative," without any discussion of the potential differences between an "image" of the motor output and the "image" resulting from sensory input. In principle, for the example of the fly, it would be necessary for the fly to model the exact dynamics of how muscle commands affect wing movements, how wing movements turn the body, how the turning body affects the visual input, and how the visual sensors will respond to this input, to be able to produce a prediction that is spatio-temporally accurate. Indeed, as the response of its visual system to image velocity is known to be highly dependent on the contrast and spatial frequency of the scene, the fly must in theory have advance knowledge of the scene it will encounter as it turns if it is to calculate an exact prediction. It seems unlikely that all this is actually occurring in the small brain of the fly.

There are several ways out of this problem. The first is the recognition that the emulation may not be complete and accurate (briefly referred to by Grush in sect. 5.2), and that it does not need to be so to be useful. A simple scaling of the strength of the turn signal might be a good enough approximation of the size of the resulting optomotor signal for the fly. An issue that then arises is the following: How "schematic" (to use Grush's term) can the emulator be yet still count as an emulator? If it is defined by its role in the control cycle, as pictured in Figure 4 of the target article, then arguably even a simple inhibitory signal sent from motor areas to sensory areas should count, such as the "saccadic suppression" found in locusts (Zaretsky & Rowell 1979). On the other hand, it does not seem warranted to say that the animal is using an "internal model" in such a case. But what additional detail in the emulation would justify the ascription of a model? What if the prediction is also direction-specific, as for the *Drosophila* optomotor response (Heisenberg & Wolf 1984)? Or can be retuned online to be proportional in size (Mohl 1988)? Or is well matched in its time course, as in the inhibitory signal anticipating self-produced sound bursts in the cricket (Poulet & Hedwig 2003)? Must it estimate more than one dimension of the input signal? Or is the defining characteristic of a true internal model that it can also be run off-line for planning? One might think this would eliminate examples from "small-brained" animals, although there is evidence for route-planning in jumping spiders (Tarsitano & Andrew 1999).

A second way out of the problem of having to emulate all the details of the motor system, environment, and sensors is discussed in more detail by Grush in section 5.1. The idea is to have the emulator predict higher-level representations (such as the new layout of objects in the world relative to the observer after a particular movement), rather than the raw sensations this new layout would produce. In the case of the fly, the prediction might be supposed to correspond to the output of the higher-level neurons that integrate motion across the visual field, rather than the response of each elementary motion detector. Of course, this shortcut to avoid modelling peripheral mechanisms can apply on the motor output side as well: The emulator can take as input a high-level motor command such as "rotate ten degrees left" and assume that this

results in a ten-degree-left rotation of the body position, rather than try to model the details of the muscle movements and resulting forces that cause the turn. However, taking this approach to its logical conclusion, the “highest level” of motor commands might be described in terms of desired goals and the “highest level” of perception might be confirmation that the goals have or have not been achieved – in which case the “emulation” required for prediction becomes trivial: It consists simply of making a copy of the goals.

Thus, in both cases – allowing approximate, or high-level, emulation – we are led to consider a continuum, from complete and detailed internal models of the plant, environment, and the measurement process, to extremely simple processes that might suffice for prediction. Much of the evidence cited to support the emulation theory does not distinguish where on this continuum the actual mechanisms might lie (Karniel 2002). Our intuitions tend to be very different when discussing human brains and insect brains, yet the problems are often common ones. The force of the theory – as a predictor of future behavioural or neurophysiological findings – lies in the assumption that the mechanisms, at least for human brains, lie towards the complex end of the continuum. The interesting question is whether there is strong evidence to support this assumption.

Two distinctions concerning emulators

Mark Wexler

LPPA/CNRS, Collège de France, 75005 Paris, France.
wexler@ccr.jussieu.fr <http://wexler.free.fr/>

Abstract: The target article distinguishes between modal and amodal emulators (the former predict future sensory states from current sensory states and motor actions, the latter operate on more abstract descriptions of the environment), and motor and environment emulators (the former predict the results of one’s own actions, the latter predict all changes in the environment). I question the applicability of modal emulators, and the generalization to environment emulators.

Grush postulates two types of emulators, modal and amodal. Modal emulators operate at the sensory surface, predicting future sensory states (in some given modality) from current sensory states and motor commands; amodal emulators, on the other hand, operate on deeper descriptions of environmental parameters, predicting future parameters from current parameters and motor commands. Each of these two types of emulator runs into at least one major difficulty.

The difficulty of the modal emulator (of which Mel’s models are examples) might be called *sensory aliasing*: Most of the time, current sensory states and motor commands can lead to many (usually infinitely many) future sensory states. Although Grush acknowledges this problem, it is probably worse than he lets on. Let us start by examining two special cases in vision where sensory aliasing *might not* be a problem.

When manually rotating a solid object, if the axis of rotation exactly corresponds to the line of sight, knowledge of the current two-dimensional image and of the motor command to the hand (assuming that this predicts the hand’s angle of rotation) is enough to predict future sensory states – they are just rotations of the two-dimensional image by the given angle. This is the kind of rotation studied by Wexler et al. (1998).

In the human eye, optic and geometric centers nearly coincide. As a consequence, when the eye rotates in the head (during a saccade, for example), the two-dimensional optic array simply rotates on the retina, a rotation equal and opposite to that of the eye. Therefore, the sensory effects of eye movements can also be predicted, to some extent, by a modal emulator. (This is only true to the extent that we ignore the visual periphery, and changes in resolution and color due to a non-uniform distribution of photore-

ceptors.) This is the sort of prediction that Duhamel et al. (1992) found in monkey parietal cortex.

Now, these are just special cases of object manipulation and self-motion. As soon as we generalize even slightly, we come up against sensory aliasing.

When an object is rotated about any axis other than the line of sight, its image on the retina deforms because of the laws of optical projection, and undergoes other changes because previously occluded parts come into view. These deformations cannot be predicted from motor commands coupled with sensory data alone; knowledge of the three-dimensional structure of the object (the deep parameters) is required. The relationship between manual rotations in depth and corresponding mental rotations was studied by Wohlschläger and Wohlschläger (1998).

As soon as the head moves (and it does so nearly all the time, unless one’s teeth are clenched around a bite bar), the eyes no longer just rotate, but undergo translations as well. This leads to motion parallax, and again the changes on the retina can be predicted only if one knows the three-dimensional layout of the environment. (Interestingly, motor action seems to play a special role in this process – see Wexler et al. 2001.) Hence, Grush’s “white cube” example in section 4.4 is incomplete: Moving sideways while facing a cube leads not only to a shift of the retinal image (as would be the case for an eye rotation), but also to a deformation, a deformation that can only be predicted if one knows the three-dimensional shape of the object (a solid cube, rather than an infinite set of other objects compatible with the initial sensory data).

So situations in which modal prediction is even possible appear, at best, as special cases of much more widespread situations in which sensory aliasing precludes any possibility of predicting sensory states from other sensory states. It would be strange, therefore, if a modal prediction mechanism existed just for those special cases, only to be superseded by a much more general amodal predictor as soon as one steps out of the special case. By parsimony, we might suppose that all predictors are amodal.

Another significant distinction to make is between motor and environment emulators. Grush begins with motor emulators – ones that predict, modally or amodally, the results of one’s own action – but winds up with something he calls “environment emulators,” the much more general task of which is to predict changes in the environment that are due not only to the subject’s own action, but to all relevant forces. Now, it might be argued that motor prediction is very special, in that one has intimate, before-the-fact knowledge of forces and goals. Indeed, the dynamic properties of the predicted trajectories of objects that one manipulates seem to be quite different from those of objects one merely observes (Wexler & Klam 2001).

The main difficulty is how the environment emulator would articulate with the motor emulators described in the first part of the target article. Indeed, it is not easy to see how the precise and well-documented mechanisms of forward models, Kalman filters, and so on, are to apply to the environment emulator. What does the environment emulator share with the motor emulator, other than the general concept of prediction? The answer to this question might prove quite interesting. For instance, it *could* turn out that to predict and to imagine changes in external objects independent of one’s own action (“motion-encoded” prediction, to use Kosslyn’s term), one nevertheless has to “imagine” acting on the objects on some level. In this case, the activation of the motor system would not rise to the level of consciousness, but would result in the activation of the motor predictor. However, without such a motor implication in the prediction and imagination of external events, it is hard to see how one can make the leap from motor to environment emulators.

Computing the motor-sensor map

Oswald Wiener and Thomas Raab

Kunstakademie Düsseldorf, Eiskellerstrasse 1, D-40213 Düsseldorf,
Germany. postmaster@kunstakademie-duesseldorf.de tr@sil.at

Abstract: “Articulate models” subservient to formal intelligence are imagined to be heterarchies of automata capable of performing the “symbolic (quasi-spatial) syntheses” of Luria (1973), where “quasi-spatial” points to the abstract core of spatiality: the symbol productions, combinations, and substitutions of algebraic reckoning. The alleged cognitive role of internal “topographic images” and of “efference copies” is confronted with this background and denied.

An organism receives stimuli of a particular character and emits regular responses that are in general of an entirely different character. It can be regarded as computing a mapping from the afferent domain to the efferent range. Vice versa, the organism’s environment “computes” a mapping from the organism’s motor (etc.) output to a range of physical configurations; a subset thereof feeds back to the organism as updated stimuli. It seems useful to distinguish between a category of organisms that can muster simulations of the environment’s computations in order to aid their own transformations of sensory input into motor commands, and a broader category of organisms that compute their responses more directly.

Let us for historical reasons call a mechanism capable of computing a particular mapping of the latter kind, a “sensorimotor schema.” A particular sensorimotor schema does not have parts that co-operate in analogy to co-operating parts of the environment. It is a spatiotemporal mould for a particular class of stimuli. It cannot properly be regarded as a representation of that class (otherwise a hand, for example, would have to be termed a representation of the class of objects that it handles). Organisms relying exclusively on sensorimotor schemata exhibit various kinds of action readiness; they cannot correctly be said to entertain expectations. They remain on the level of what Grush has termed “sensation.”

Organisms of the first kind, however, internalise relevant regularities of their environment by creating a realm of “formal schemata” that sits, as it were, on top of the realm of sensorimotor mapping and dominates it. Formal schemata work online by taking prerogative over input to the sensorimotor realm in order to establish perception proper, or by controlling sensorimotor control of efference. They can work offline to predict the outcome of external action and of state transitions of other schemata. They are necessary for the construction of new formal schemata qua solutions to problems, or can be used to supervise the building of new sensorimotor schemata to mitigate demand on the resources necessary for running formal schemata.

Furthermore, formal schemata can sometimes be used to infer generically what a particular external setup would look like, which amounts to problem-solving in the visual domain (the solution is a set of operations that, when executed, produce a picture in the environment). Because perception is the control of models by stimuli originating in external objects, the operations of the inference will produce some of the concomitants of seeing by priming sensorimotor schemata that would react to the relevant visual signals, should any of their kind arrive, and by pointing to formal schemata that could be activated in order to include additional information. This setting up of pointers is not a measurement (unless one wishes to treat “effect” and measurement as synonyms).

Whereas a sensorimotor schema can be conceived of as an “analogue” machine, a formal schema may didactically be compared to a string-processing device – a sequential, discrete-state, discrete-symbol automaton epitomised in the notion of the Turing machine (Wiener 1998; with a view to neuropsychology, a formal schema should rather be construed as a layered Petri net; with a view to introspection and “mental images,” as a formation of Moore automata). Embedded in suitable running environments (heterarchies of other formal as well as of sensorimotor schemata that provide qualified control and parameter signals mediating, where

applicable, connections with the external environment), any particular formal schema in an organism is indeed what seems to be intended by the term representation (utilisable *as such* for the organism): a *model* (Wiener 1988), a *structure* (Wiener 2002) of (the set of stimulus sequences that for the organism constitute) any object belonging to a particular set.

The simulations of the behaviour of external objects or even of the organism’s own sensorimotor schemata (Piaget 1947; Rozin 1976), as accomplished by formal schemata, rest on functional equivalences of analogue devices and suitable symbol-processing automata. In their interactions the models generate and accept symbols that stand in bare *functional analogy* to stimuli on the one hand, and, on the other, to the organism’s own efferent signals. An activated model’s offline behaviour is thus determined by its own spatiotemporal regularities, as parameterised by the actual running environment and by other models’ output functionally equivalent to the end products of sensory processing or to motor commands. It seems useful to assume the same format for all model output – no modal buffers, no efference copies effective on the model level. In the last analysis, mental images are pointers (to operations, viz., to models) that for their part can be operated on. Readers interested in the notions given are referred to Wiener (1996; 2000, and forthcoming).

Look again at Figure 7 of the target article. Because of some motor command, the effector changes the environment. The outcome of this event is measured, and the result of the visual measurement is a retinal image (sect. 4.4). The same motor command changes the state of the image emulator, which, as an emulator of the sensor sheet (sect. 4.4), now carries another fresh quasi-retinal image. At the same time, the motor command changes the state of the organism/environment model, and a modal measurement of the result projects a third image onto the inner eye’s retina. The three images are somehow compared, and the resulting sensory residual – perhaps a pixel array having all the pixels active that did not partake in the formation of every single picture? – is used to correct the picture in the image emulator – perhaps by being added algebraically? – and to rectify the model state by virtue of a mapping of the sensory residual to, presumably, differential motor commands. Feedback from the model to the controller (the “motor centers”) can now proceed: that is, information about the objects and states in its egocentric environment (sect. 5.1). The controller will somehow compare this information to its goal, itself specified in terms of objects and states in the environment (sect. 5.1), in order to update its motor commands.

Grush embraces the idea that perception is contingent upon matching the shapes of “percepts” and of internally generated images (Kosslyn & Sussman 1995). The principal component of his image emulator is a Kosslynian “visual buffer,” implemented somehow in the manner of Mel (see Mel 1991). But Mel’s dealing in changeable pictures, with their pixel transformations governed by neighbourhood relations, is successful – so far as it goes – only because there is just one moving complex (the “arm”) that is accordingly identified without trouble. At the same time, this appears to be the most complicated case solvable by devices of that kind (we will not speak of Mel’s run-of-the-mill backtracking mechanism that benefits from “representation without intelligence”). If such two-dimensional arrays of connectionist pixels were to carry geometric projections from spaces of higher dimensionality, and if the objects in those spaces were flexible and moving around independently, correct pictorial transformations and object recognition and tracking (even the “recognition” performed by contemporary computers) would be impossible. This consideration renders Grush’s image emulator and the corresponding part of the target article’s Figure 7 gratuitous.

Similar problems derive from Grush’s adaptation of von Holst’s and Mittelstaedt’s notion of “efference copy” (cf. von Holst 1954). With the original authors the concept works because, again, the conditions they assume are quite clear-cut: The efference is oculomotor only, and the task concerns the entirety of the visual stimulus complex at the time of a single saccade. But to compute a

mapping from motor commands to pictures, a device would need input unique to each picture, a condition certainly not met by motor “plans” for human effectors. Of course, the lay would be different if the device were a model of the pictured object, already activated by other signals and thus all by itself constituting a specific interpretation of the accepted signals. We acknowledge that Grush points to this possibility, but in the case now at hand the input could be anything versatile enough to exploit the capabilities of the respective model.

Another problem of the same kind arises concerning the control of the controller. If Grush’s sensory residual is questionable because the idea of picture subtraction does not make sense cognitively (while the idea of structure tests does), then the comparison of “environmentally specified” goal signals to feedback in the same format will work even less – unless the controller is able to use the model realm or, rather, reduces to a servomechanism of the latter. That would render the idea of model control by efference copies gratuitous.

Our arguments imply, or so we believe, that the idea of measuring forward models and Kalman filtering the results, if at all applicable to intelligent organisms, applies only to functions well beneath and possibly up to the level of sensorimotor schemata.

Motoric emulation may contribute to perceiving imitable stimuli

Margaret Wilson

Department of Psychology, University of California at Santa Cruz, Santa Cruz CA 95064. mlwilson@cats.ucsc.edu
<http://psych.ucsc.edu/Faculty/mWilson.shtml>

Abstract: First, I note three questions that need further exploration: how fast the emulator operates, compared to the real-time events it models; what exactly perceptual emulation, with no motor component, consists of; and whether images are equivalent to raw sensations. Next, I propose that Grush’s framework can explain the role of motor activation in processing “imitable” stimuli.

A few questions . . . I shall begin by noting a few problems, puzzles, and gaps in the model. These are not intended to detract from the considerable accomplishment of Grush’s proposal, but rather to suggest lines along which this set of ideas could be pushed forward.

(1) *The timing problem.* How closely in sync with the plant (the system being modeled) is the emulator? On the one hand, an emulator is supposed to be fast, bypassing delays in sensory feedback and allowing corrections in motor control to be made in a timely fashion. This requires an emulator that runs its simulation almost simultaneously with (perhaps even slightly ahead of?) the events in the external world. On the other hand, however, an emulator is supposed to correct for measurement error, by running measurement processes in reverse and employing Kalman filters. This is an architecture that adds several processing stages beyond simply sending sensory feedback to the controller, hence actually exacerbating the first problem that the emulator is asked to solve. Can both these functions be served by the same emulator, or is this architecture being asked to do too much?

(2) *The perceptual simulation problem.* Grush acknowledges that much of human cognition resides within the boxes in his Figure 7, rather than in the functional relations between them. In some respects, this is completely justified. For example, the box modestly labeled “measurement inverse” actually represents all of perceptual processing (to wit, reconstructing the distal stimulus from the proximal stimulus); but providing an account of this box is not the purpose of Grush’s project. More problematic, though, is the box labeled “organism/environment model.” It is here that, in the “degenerate case” of perceptual imagery without a motoric component, the entire business of imagery happens. And explaining im-

agery is a central part of Grush’s purpose. Somehow, the “organism/environment model” has internalized principles and regularities that allow it to run simulations of the external world. How such a feat of simulation is possible seems to be the real heart of the imagery question. (Some possible beginnings of an answer to this are suggested below.)

(3) *The images-aren’t-pictures problem.* Grush takes at face value the idea that modality-specific imagery is equivalent to unprocessed sensory input. In fact, this idea has been vigorously challenged (e.g., Reisberg & Chambers 1991; Reisberg et al. 1989). If a modality-specific image at all resembles a percept, as opposed to a raw sensory input, then the model in Figure 7 would need to be changed. It would not be possible to generate this kind of imagery by going all the way to an amodal three-dimensional representation (the organism/environment model) and then simply “looking at it” (the second measurement) to recreate the raw sensory input.

. . . and a proposal. Having issued these challenges, though, I would now like to turn to a different area in which Grush’s model may offer considerable explanatory power: the perception and representation of stimuli that can be imitated with one’s own body. There is growing evidence that such stimuli (which consist primarily of other humans’ postures and movements) have a special status in the human cognitive system. Specifically, imitable stimuli appear to have privileged connections to motor representations of performing similar postures and movements with one’s own body. Evidence for this comes from the literature on mirror neurons and on stimulus response compatibility, and from a variety of other sources (see Wilson 2001 for review). What is the purpose, though, of activating imitative motor programs in response to perceptual input, when there is no intention to overtly imitate? One possibility is that motor representations may actually play a role in perceptual processing.

As Grush suggests, an emulator which can capture predictable properties of perceptual events and simulate them in real time may serve both to assist perceptual processing and to recreate perceptual events in their absence. Indeed, the operation of just such a system appears to be behind the phenomenon of representational momentum, in which perceptual events are immediately misremembered as having continued beyond the point at which they ceased to be visible.

In most cases, representational momentum occurs for movement trajectories that are predictable based on simple geometric principles, such as rotation about an axis (Freyd & Finke 1984), circular paths (Hubbard 1996), spiral paths (Freyd & Jones 1994), oscillatory motion (Verfaillie & d’Ydewalle 1991), and predictable changes in direction such as bouncing off a wall (Hubbard & Bharucha 1988). These are the kinds of predictions one could build into a perceptual emulator without much difficulty (see Question 2, above). However, representational momentum also occurs for complex human motion (Verfaillie & Daems 2002; Verfaillie et al. 1994).

Based on what principles is the emulator able to predict the trajectories of complex biological motion? A likely answer is: based on principles of body biomechanics. That is, the emulator may have internalized the constraints of human body movement – the range of motion of joints, acceleration properties of muscles, and so on. In Grush’s terms, the emulator may be an articulated model, isomorphic in its parameters to the relevant parameters of the human body.

From here it is only a small step to suggest that activation of motor brain areas in response to perception of conspecifics may be a functional part of an emulator, assisting in the perceptual prediction of those events. That is, just as in the case of emulating motor events and the case of emulating perceptual events as they are altered by self-motion, emulation of imitable stimuli may receive part of its drive from the controller. In each of these cases, the controller issues motor “commands” in the service of representation rather than (or in addition to) in the service of overt movement. The perception of conspecifics, then, may form an additional class of cognitive events that the emulation theory can help to explain.

Author's Response

Further explorations of the empirical and theoretical aspects of the emulation theory

Rick Grush

Department of Philosophy, University of California, San Diego, La Jolla, CA
92093-0119 rick@mind.ucsd.edu http://mind.ucsd.edu

Abstract: The emulation theory of representation articulated in the target article is further explained and explored in this response to commentaries. Major topics include: the irrelevance of equilibrium-point and related models of motor control to the theory; clarification of the particular sense of “representation” which the emulation theory of representation is an account of; the relation between the emulation framework and Kalman filtering; and addressing the empirical data considered to be in conflict with the emulation theory. In addition, I discuss the further empirical support for the emulation theory provided by some commentators, as well as a number of suggested theoretical applications.

Although space considerations have prevented me from addressing all of the issues raised in the commentaries, I have tried to respond to those that struck me as the most important, or for which some sort of response seemed most called for. I have tried to impose some order by organizing the issues and my replies under ten headings. For those interested in particular commentaries, I have listed under each heading the commentators mentioned in that section.

R1. Kalman filters are a perspicuous special case of the emulation framework, not identical to that framework [Donchin & Raz; Goussev; Merfeld; Stojanov & Bickhard]

A number of commentaries questioned the applicability of Kalman filters (KFs), narrowly construed to various applications addressed in the target article. But the emulation framework is not identical to Kalman filtering. KFs are a specific instance of the emulation framework, but so are, for example, Smith predictors. I will try to explain below what is, and what is not, essential to the emulation framework as I intend it.

The emulation framework is (as described in sect. 2.4 of the target article) an information processing framework in which a system constructs and maintains an emulator of some domain with which another component of the system, a controller, interacts. The controller component uses this emulator in order to do at least some of the following: (1) help overcome feedback delays in online interaction with the represented domain by operating the emulator in parallel with that domain and using its feedback rather than delayed feedback; (2) run the emulator in parallel with the represented domain, even in cases where there is no feedback delay, in order to process sensory information intelligently; (3) run the emulator in parallel with the represented domain in order to form expectations that can be of use in sensory processing (e.g., anticipating where edges will be so that early visual systems can begin some processes earlier than would otherwise be possible); (4) run the emulator off-line in order to see what a certain course of action might lead to (planning), or to train the

controller (imagined rehearsal to improve skills), or just for fun (dreaming).

In introducing this information processing framework, I chose to illustrate it with the Kalman Filter (KF). The KF is an example of the emulation framework, as described in the previous paragraph. Furthermore, it is a particularly perspicuous example in that it can be described completely and rigorously without too much mathematical formalism, but when described it is clear enough to perspicuously exhibit the different components of the emulation framework (emulators, measurement, driving forces vs. process noise, etc.) that would be difficult to introduce clearly without some sort of more or less formalized model in hand. (As an example of the kind of lack of clarity we get with merely qualitative explanations, see the confusions surrounding the “simulation theory” of motor imagery that I discussed in sect. 3 of the target article.)

Nevertheless, the emulation framework is not identical to Kalman filtering (this is the answer to **Stojanov & Bickhard's** final question). I did mention this a few times (see especially the final two paragraphs of sect. 2.4). Perhaps **Goussev's** suggestion that, “We could expect the appearance of a more general nonlinear theory which will be able to embed the Kalman filter theory, likely as it did with the Wiener filter theory,” might be right. Certainly Goussev's suggestion that bare Kalman filtering has limitations is one I fully agree with.

Merfeld appears to have understood the proposal entirely correctly as very similar to the *observer framework* (see Merfeld's commentary for references); Kalman filtering is an *optimal* observer framework. And **Donchin & Raz's** suggestion that, for some applications, gating (as described by these commentators) would have more obvious relevance rather than the filtering mechanisms which are required by the KF, is correct, and in line with the emulation theory.

R2. The emulation framework is not an attempt to explain everything, nor does it always posit a single emulator [Charles; Donchin & Raz; Calvo Garzón; Hubbard & Ramachandran; Sadato & Naito; Walter; Wang & Yang]

There are two clarifications about the emulation framework that need to be made. The first is that the emulation framework does not posit that there is *a (single)* emulator in the brain somewhere which is responsible for all the phenomena discussed in the target article. Rather, it posits that the CNS makes use of an information processing strategy. Thus, there may be many emulators in the CNS, some of them related – as perhaps modal and amodal emulators of the same domain are related – as described in section 4.4 of the target article; others may not be related at all, being driven by different kinds of signals from different areas, and geared to the emulation of different domains. For example, the emulator of the musculoskeletal system used to aid in motor control is probably completely separate from the emulator implemented in the posterior parietal cortex (PPC) to anticipate the results of saccades. So, for example, when **Donchin & Raz** remark:

In his conceptual framework, Grush argues that modeling is a common theme in activities that involve fashioning our own behavior, predicting the behavior of others (i.e., theory of mind),

or expecting changes in the environment. Grush implies that this general network manifests in converging neurophysiological mechanisms.

this was not something I intended to imply. Now, it may be that some functions not only rely on the strategy of emulation, but may involve some of the same emulators (for example, it was argued that visual imagery and visual perceptual processing exploited at least some of the same emulator[s]). But in other cases the emulators may be entirely distinct with regard to the neural substrates, source of corollary discharge, and system represented.

The same assumption seems to be made by **Calvo Garzón**, who says that “Grush favors an *articulated* reading of emulation” (emphasis his). For some applications I think that the emulators involved are articulated – for example, an environment emulator, in which I don’t merely represent some big undifferentiated scene, but rather, represent parts as distinct components of the scene. (See also the commentary by **Sadato & Naito** which presents data for an articulated musculo-skeletal system [MSS] emulator.) But for others they may not be. **Charles** points out that some things that might count as process noise when only the body is taken to be the process, become predictable if the things outside the body that are causing disturbances can be known. This is correct, and can be addressed by recognizing that, in addition to the MSS emulator, there are environment emulators that do exactly what Charles suggests: they model the antics of the environment and how the body will interact with it. A wonderful example is **Hubbard & Ramachandran’s** discussion of the size-weight illusion, which results, as they hypothesize, from the interaction of these emulations (this is discussed further in sect. R6).

The second clarification involves what is perhaps an instance of a bad assumption unreflectively employed by many researchers in cognitive neuroscience. The emulation framework does not claim that it explains *all* the functions that the CNS executes. More crucially, in the case of a function for which the CNS does use emulators, it is not assumed that the CNS must *always* employ emulators for that function, or for all aspects of that function, or even that the CNS uses one scheme exclusively for any single function. Take motor control as an example. It is very likely that the CNS employs a number of different control or information processing schemes for motor control. It is likely that chewing and walking *usually* (see below) make use of central pattern generators governed by reflex arcs and simple modulatory central commands, with little or no involvement of anything like emulators or micromanaging of the temporal profile of muscle tensions. Other motor actions, such as quickly pointing to a star, probably only minimally involve reflex arcs, and probably have a lot of central involvement of emulators and micromanaging of, for example, kinematics. And even for a given function like chewing, it might be the case that sometimes, such as when I am chewing gum absentmindedly, one sort of scheme is being used more or less exclusively; and when I am deliberately chewing something, paying attention to the exact movement of my jaw – perhaps because I have a very sore tooth – that control is much more micromanaged, maybe involving emulators and with very little control entrusted to central pattern generators and reflex arcs.

Thus, the answer to one of **Charles’s** questions is: no, emulators are not always needed for all motor control tasks. This also partially responds to some of **Walter’s** worries

about the capacities of the emulation framework to handle complicated dynamics. Some kinds of motor control, such as maintaining out-of-phase limb motion, may not involve emulators or other sophisticated mechanisms at all, but may involve simple pattern generators.

I say that this is perhaps an instance of a more widespread illicit assumption because a significant factor in a number of arguments taken to support this or that theoretical framework in various domains involves just such an assumption: evidence in favor of the hypothesis that the brain does *X* is taken to be evidence against the hypothesis that the brain does *Y*, when in fact it is not, unless it is also shown that there is evidence that the brain does *X exclusively*. The brain is incredibly complicated, and has a very long evolutionary history peppered with idiosyncratic pressures and solutions. We simply cannot assume that there is *one way, one control structure, or one information processing strategy* that the brain uses exclusively even for the same task, let alone for different albeit related tasks. In particular, when I argue that the brain uses the emulation framework for motor control (for example), I should not be taken to thereby imply that other control schemes, such as Feldman’s λ model (see sect. R3 of this response), or bare closed-loop control schemes, and so on, have no application. And, conversely, evidence to the effect that the CNS uses closed-loop control, or whatever, does not constitute evidence against anything I want to claim, unless it is somehow shown that the framework in question is exclusively used for such functions (and is incompatible with the emulation framework; see sect. R3).

Motor control will be discussed in more detail in the next section, but other commentaries seemed to turn on the same point. **Calvo Garzón** sets up a dilemma based on the advantages and disadvantages of continuous and discontinuous mechanisms of emulation. But a dilemma arises only if just one emulatory mechanism is allowed. There may, in fact, be a number of mechanisms, some that Calvo Garzón would classify as continuous, and others, as discontinuous. **Wang & Yang**, for example, point out that certain logical operations cannot be explained by emulators. Not only do I agree with this, I pointed it out myself in the second paragraph of section 6.2 where I explicitly stated that the emulators themselves, which are posited to be the individual mental models, must be “manipulated by a system capable of drawing deductive and inductive inferences from them.” I am happy to add quantification and other operations to this list. Such operations are not emulators in themselves, nor are they part of the emulation framework. But the hypothesis is that these functions operate over representational structures that include mental models embodied in emulators.

R3. The emulation framework is not in conflict with equilibrium point models, λ -models, or embodied dynamical models of motor control [Balasubramaniam; Charles; Latash & Feldman; Walter]

The fact that I introduced control theoretic concepts by means of motor control examples that used joint torques, or angles, and so on, as the control signals, apparently caused some confusion. In particular, it gave some commentators the impression that the emulation framework involves a

particular stand on the specific details of the control signal and feedback. Notice that my characterization of the emulation framework in section R1 here did not make any mention of the specific nature of the commands or feedback. It doesn't need to.

The problem is partly historical. Early motor control research often assumed that the brain's commands were in terms of joint torques, angles, and so forth. We might call such schemes *robot control*, in the spirit of the title of **Latash & Feldman's** commentary. Because I used examples that were robot control-ish in nature to introduce the control theoretic concepts, some commentators got the impression that I was endorsing robot control over one of these alternatives (which I will discuss shortly). But the emulation framework is not a stand on these specifics. It is a stand on different specifics. Adding to the confusion is that, historically, proponents of robot control have been among those who argued for the employment of forward models, since feedback delays can cause stability problems for robot control schemes. That this might be a factor causing confusion is suggested by the following remark by Latash & Feldman:

The idea of movement production by shifts in equilibrium states avoids many of the problems that emerge when control schemes are borrowed from an area of technology (e.g., robotics) where movements are powered by predictable actuators, not variable, spring-like muscles. Such control does not need an on-line emulator to lead to stable behavior.

But the target article didn't mention stability *at all*, let alone as a motivation for forward models or emulators. Rather, feedback delays (which *can* raise stability issues, but can cause problems apart from these issues), sensory surrogacy, and producing better feedback estimates were among the main motivations. And crucially, *these motivations are in full force for schemes other than robot control, such as equilibrium-point models, and Feldman's λ model.*

In Feldman's λ model (not identical with but similar to equilibrium point models; and, apparently, the favored model of both **Balasubramaniam** and **Latash & Feldman**), the CNS implements a hierarchical control scheme in which a superordinate controller sets one or more parameters that influence the operation of one or more subordinate controllers. The variable controlled by the superordinate controller is effectively a bias that influences the muscle length recruitment threshold of motoneurons. (Briefly, motoneurons get inputs from stretch receptors such that the length of a muscle will have an effect on when a given motoneuron will become active. We can call the length to which a muscle has to be stretched, to activate a given motoneuron, " λ " – hence the name of the model.) But setting this parameter does not determine exactly what the body will do. Rather, a given parameter setting, together with external influences and perhaps the operation of reflex arcs, determines what the body does. (For more detail, see Feldman & Levin 1993; 1995. The dynamics of such low-level systems is part of **Charles'** concern.) In fact, to a large degree, I'm a fan of this and related models, and it may well be that a number of the complicated phenomena **Walter** points to are handled by such schemes.

Note that in many of the cases discussed by Feldman and collaborators, the superordinate controller is open loop: On the basis of some goal (like *stand still*, or *walk*, or *hold the weight*) this controller issues a command in the form of a pattern of λ settings, and then subordinate mechanisms,

themselves typically described as closed-loop systems, interact with the environment and other low-level motor machinery to produce the behavior. This is analogous to you setting the temperature on a thermostat: you are a superordinate controller that sets the knob and then walks away (you are operating "open loop"); but the knob setting sets a parameter that influences how a subordinate system, the thermostat, will interact with the environment. You don't micromanage the operation of the heater or air conditioner, you may have no idea how any of that works. And your control signals and goal-state specifications are not represented in a form even similar to those used by the thermostat. You deal with felt temperatures and knob twists, the thermostat deals with electrical voltages in wires.

But note that the superordinate controller (both you controlling the thermostat, and the entity setting λ values) might need to operate "closed-loop." After a while, the room is too warm, and so you go back and adjust the knob. You keep adjusting the knob based on your comfort level until it feels good. In this case the superordinate controller as well as the subordinate controller are closed-loop. Though again, each is dealing with feedback and commands in very different formats. Ditto for the λ model. If the first group of λ settings didn't get the job done correctly, a new one might have to be produced. With this remark we get to what is perhaps the biggest source of confusion on this topic. In some cases, the feedback needed to ensure the appropriate behavior is handled by the subordinate controller(s). And for commands like *stand still*, or *walk*, this is probably the case much of the time. Because of this, the superordinate controller need not micromanage things, or even be more than an open-loop controller. But other kinds of cases must involve feedback processed by the superordinate controller. When I decide to point to a star, whether or not my arm does the right thing is not something that can be assessed by stretch receptors and spinal reflex arcs. They don't know where the star is. Rather, the superordinate controller must get feedback (visual, in this case) and assess the extent to which the goal was met, and alter the control signal if it was not. *And this is true regardless of what format the commands are in: robot control, equilibrium point settings, λ values, whatever.* Unless λ models and equilibrium point models are magic, then for goal-directed actions the initial equilibrium point or λ setting commands might not be accurate, and adjustments might need to be made. That is, the superordinate controller might need to be closed-loop.

Now, if adjustments are made, presumably this is on the basis of sensory feedback to the effect that the initial command isn't getting the job done. If such feedback is always very fast and never in error, then it could be relied upon. But if it is even occasionally delayed, or if it is imperfect and hence could benefit from some sort of processing or filtering, then mechanisms like those posited by the emulation framework have application. The emulator would be some mechanism that, on the basis of information about the current state of the MSS and the motor command (whatever format this command is in, λ settings, equilibrium points, joint torques, whatever), produces an estimate of what the forthcoming feedback signal will be (visual, kinaesthetic, it doesn't matter – *whatever* format the feedback is in that would lead the superordinate controller to adjust its commands, can be learned as the format produced by the emulator). In any case, such a scheme might provide faster

feedback that would be of use. Furthermore, even if delays were not a problem, imperfect sensors might be, and hence having an a priori estimate to combine with the sensor signals to get more accurate feedback could still be a benefit. And even if the sensors were always fast and perfect, having the wherewithal to try out motor plans before deciding on a course of action might be of benefit. In short, all of the benefits that were argued in the target article to motivate the use of the emulation framework are, as far as I can tell, in full force even on these alternate schemes. If they are not, the case has yet to be made.

R4. The emulation framework provides a good theoretical grip on the notion of representation [Dartnall; Jordan; Newton; Stein; Stojanov & Bickhard; Wang & Yang]

An implicit contention of the target article (one explicit, unfortunately, only in the title) is that the emulation framework is not just a theory to the effect that the CNS employs emulators, but rather that, precisely *because* the CNS employs emulators, it *represents*. This idea is challenged by **Wang & Yang** and **Stojanov & Bickhard** in two very different ways. Wang & Yang raise considerations aimed at showing that emulation is not *necessary* for representation, since lots of things that don't appear to be involved with emulators do appear to be representations. Stojanov & Bickhard push in the other direction, questioning the *sufficiency* of emulators. The result is a powerful tag-team challenge.

Before entering the ring, though, a few words on what it means to have a theory of representation. *All* physical entities – planets, paramecia, cars, and Nobel Prize winners – can have their behavior explained by appeal to the states of their surroundings, their own (internal) states, and applicable laws. This is true *in principle* at the level of basic physics; and also *in principle* at the level of chemistry and perhaps biochemistry. But some entities are such that, in addition to these in-principle physical explanations, at least some of their behavior is amenable to a kind of psychological explanation: explanations that appeal, *inter alia*, to contentful states and operations over them. For example, when my hand moves out towards a glass, one could in principle explain that action purely in terms of the physical properties of my surroundings, neural and bodily states, and so on. But one could also provide a psychological explanation: I was thirsty and wanted a drink, and I believed that the glass has cool water in it. Now it may be that in many cases such psychological belief-desire explanations have limitations. They can't explain *everything* we do. Nevertheless, they sometimes (I think, quite often) work, and this is something that needs to be explained. How can a system that is governed by purely physical laws and whose behavior is explainable (in principle, if not in practice, on account of the complexity of the system) by these laws, also be one that can, at least sometimes, be explained quite well using a very different sort of theory? To my mind, the key notion here is representation. We need an account of how something purely physical can be *about* something, can carry a content. How can some neural state in my head be about the water in the glass, or about the Eiffel Tower (especially when I am nowhere near the Eiffel Tower, or the glass is actually empty, though I believe it to be full)?

Wang & Yang claim that emulation is not necessary for

representation. I think it is clear from their commentary that they have in mind a notion of “representation” very common in cognitive science: *problem solving*. Now, (1) problem solving, and (2) being amenable to psychological explanations involving contentful states, are two different things; and so it would be no surprise if an account meant to explain (2) didn't strike people interested primarily in (1) as adequate. To see that (1) and (2) are orthogonal, note that many problems can be efficiently solved without anything close to content-bearing states being involved (the Watt governor is a favorite example); and many cases of contentful contemplation have nothing to do with problem solving, and may even get in the way of problem-solving success: as when I (wrongly) believe that the water is poisoned and so fail to drink a clearly visible and potable glass of water when my body is dehydrated.

It doesn't matter whether those interested in (1) or those interested in (2) get to decide what the right use of the expression “representation” is, so long as we remain clear on what we are trying to explain. If one is interested in (1), then external props like sticking orange stickers on the buttons one needs to press on an answering machine in order to get it to play correctly is a representation, because it is something that aids problem solving. But it aids in solving this problem not by representation in the sense of (2), but by being salient, or something like that.

The other element focused on by **Wang & Yang** is externality. They appear to go from something's being external to its not being amenable to a representational analysis by the emulation theory. I have a saying with which Wang & Yang should agree: *the skull is metaphysically inert*. Being on the inside versus on the outside of a bone barrier shouldn't by itself carry any weight as to the question whether that thing is a representation (whether of type [1] or type [2]). What we need is a theory of representation. Now, if one's theory is that anything involved in problem solving is a representation, then it turns out that just about anything is or can be a representation. Which is fine, if (1) is what one is interested in.

But for those interested in (2), the emulation theory gives us a different angle on the issue. According to it, something is a representation if and only if it is used by some system to stand for something else, and the “stand for” is explained in terms of use. The motor areas use the body emulator to represent the body by driving it in the same way (by means of the same kinds of commands) that it would use to control the body; the ship's crew uses the map and marks on it to represent the ship and its location by manipulating it by processing the same kinds of commands that would also drive the ship. But the orange sticker on the correct answering machine button is not used to stand in for anything. It does its job simply by being salient and well-placed.

Now, according to the emulation theory, emulators can be realized in neural circuits, and this was discussed in a number of examples in the target article. But they can also be realized externally: for example, if I am playing chess with a friend, and use another board to try out moves so that I don't have to disturb the actual board (I discuss this example at length in Grush 1997). So the information processing framework not only does *not* claim that representations must be internal, it actually gives us some way to distinguish states that aid in problem solving but aren't *about* anything, from those that are about things, even for external entities.

Newton, who comments on the ability of the emulation proposal to provide some theoretical grip on the slippery notion of representation, understands the proposal exactly correctly. I discuss the notion of representation as essentially one part of a three-part relation – in a way similar to Newton’s discussion – in Grush (1997). I would only add that so long as their role in representational activity is understood as constitutive of their status as representations, speaking of the entities in emulators – whether internal as with visual imagery or external as with the chess board – as being representations seems correct. That is, objects can be representations. But only insofar as these objects play a role in a certain kind of process. I think that Newton and I are in agreement on this.

Now to **Stojanov & Bickhard**. Their challenge pushes in the opposite direction: is emulation *sufficient* for representation, even if we are interested in (2)? They couch the issue in terms of the normativity of representation, but for the readership who are not philosophers it will perhaps be easier to frame it as a question concerning what it is that makes emulators and their articulators (if any) represent the target domain and its parts? This is not a question only about the emulation framework, but about any account of representation (in the sense of [2]). A common assumption is that causal or informational factors can settle representational questions (and in fact this is how most neuroscientists often use the expression “representation” – the neural firings represent the things that cause them to fire, like faces or action verbs, etc.). But this assumption faces conceptual problems that I won’t recount here.

Stojanov & Bickhard point out that two other approaches to this question have difficulties. The first view is that state *S* represents content *C* if someone interprets *S* as meaning *C*. For example, some ink spots mean the Eiffel Tower (you have just seen an example), and they do so because we interpret them that way. Apart from our interpretations, they would be meaningless spots. This view has two problems: one, it appeals to an interpreter who, presumably, must be able to represent content *C*, and we thus need another interpreter to explain the first interpreter’s semantic wherewithal. And, second, it seems intuitively wrong that my mental states – my own brain states – have meaning only because someone else is interpreting them.

The second approach **Stojanov & Bickhard** don’t like is the view that (for evolved nervous systems), a nervous system state *S* represents content *C* if (and only if) *S*’s carrying information to the effect that *C* explains why the neural structures that support *S* were replicated. I share Stojanov & Bickhard’s scepticism about both of these approaches, and for probably the same reason: Whatever explains representational efficacy must be state-determined. Whether or not my brain is actually representing what it represents, cannot be up to external interpreters or to things that happened hundreds of millions of years ago. If nobody was interpreting me, I would still be representing my environment. And if for some miraculous reason it turns out that I was constructed in a lab rather than born (hence lacking in any evolutionary history), I would still be representing my environment. And the explanation would be in terms of things going on *in my brain now*. (I have discussed these topics in more depth in Grush 2001.)

So why do the emulators or their articulators represent? One might think that, to the extent that the emulation theory is trying to give an account of genuine representation,

it falls into the familiar homunculus fallacy: the emulator is a representation because it is used by the controller to represent the target system. But then doesn’t that make the controller a little homunculus? And doesn’t this homunculus need its own account of representation? And then aren’t we off on an infinite regress just like the “interpretation” theory discussed above?

According to the emulation theory, the emulator represents the target domain not because the controller interprets it or its states to be about the target domain. Interpretation is itself a semantic notion, and so we would need an account of how the controller manages to represent, and we have the regress. Rather, the emulator and its states represent because the controller *uses* it and its states as stand-ins for the target system. And *use* here can be understood perfectly naturalistically, that is, non-normatively, in terms of dynamical coupling or whatever, and does not require any appeal to a semantically question-begging notion of interpretation (cf. Grush 1997).

In addition to the double-team by **Wang & Yang** and **Stojanov & Bickhard**, the emulation theory’s attempt to provide some insight into the notion of representation is challenged at an even more fundamental level by **Jordan**, who thinks that the notion of representation itself needs to be abandoned. Jordan says that my

model begs representationalism because he [Grush] begins by conceptually dividing the problem into organisms and environments. Given this dualism, the task becomes one of determining how it is that organisms build models of the environment in their brains in order to get around in the world. This then sets the stage for the introduction of yet another dualism – efference and afference.

First, it is not clear that the starting point of the emulation theory is a division between organisms and environments. I start with the problem of representation (2), and the account eventually constructed is transparent to organism boundaries, as the above remarks on external representations should make clear. As a contingent fact, I think that most emulators are implemented in organisms’ brains in their skulls, but this is not a fundamental assumption of the theory, and in fact the theory tells us exactly what representations are in a way that makes no reference to organisms at all.

But **Jordan**’s deeper objection is here:

At every point in this phylogenetic bootstrapping process, regularities and their control are the issue at hand. Seen in this light, Gibson’s (1979/1986) notion of “resonance,” as opposed to representation, takes on new meaning. An organism’s nervous system resonates to environmental regularities because the nervous system itself is an embodiment of those regularities.

Perhaps the sort of operation envisioned here is the entirely kosher closed-loop control, just described in different terms: specifically, in terms of two entities engaging in mutually adjusting dynamical coupling so as to produce stable behavior. That’s fine. As I mentioned in section R2, the emulation theory does not claim that all aspects of neural or cognitive function are explained in its terms. There is no doubt a great deal of closed-loop activity, and perhaps some of it is such that the above description is accurate.

But if “resonance” means simply adaptive interaction, à la closed-loop control, then it has limitations. I can think of the Eiffel Tower, of the Big Bang, and plenty else that I am not in any way resonating with in my environment (things

not even in my light cone); I can reach for the glass because I am thirsty even though it is empty, if I believe it is full. Environment-driven resonances have their limits as far as explaining the full range of cognitive tricks goes.

Finally, I'd like to turn to **Dartnall**, who draws some connections between the emulation theory and the "extended mind" view recently discussed by Clark and Chalmers (1998). Two points are pertinent here. First, as discussed above, the emulation theory does not mention skulls, and so the idea that cognition might extend beyond the skull is not a problem for the theory. I should, however, mention in passing that the arguments for the "extended mind" position by and large depend on a certain understanding of the mind – specifically as a problem solver, and hence as something that uses representations in the sense of (1). The arguments typically proceed by showing how external props are used by people to solve problems – props that are sometimes internal (e.g., pieces of paper can aid memory, which is prototypically thought of as internal, etc.). So, while the emulation theory can accommodate external representations, it does so in a way that doesn't let anything go. Not everything that aids in problem solving is a representation in the sense of (2), but only those things that play the right kind of role in the right kind of information-processing structure. But again, who gets to plant their flag in the word "representation" is not important. What is important is that we try to distinguish the different kinds of states and processes involved in the sort of phenomena we are trying to explain. Problem solving is a perfectly legitimate thing to study.

But **Dartnall's** suggestion is slightly different: it is, that the world can leak into the mind. I agree entirely with this suggestion, and in fact in Grush (2003, sect. 6) I discuss this briefly. The basic idea is that emulators in the brain are typically, if not always, constructed and maintained as a function of observing overt interaction; their ability to represent the target system is in some strong sense dependent on the target system itself, and on the details of the organism's (or other entity's) interaction with it.

R5. Empirical data potentially in conflict with the emulation framework [Gaveau, Desmurget & Baraduc (Gaveau et al.); Reed, Grubb & Winkelman (Reed et al.); Smith & Gilchrist; Tomasino, Corradi-Dell'Acqua, Tessari, Spiezio & Rumiati (Tomasino et al.)]

Smith & Gilchrist discuss some nice results from experiments in which subjects showed a bias in the location from which they begin drawing a triangle as a function of prior imagery. The authors take this to be unpredicted by the emulation theory:

It seems that mental imagery is not an entirely off-line process. . . . Grush tacitly takes the position that mental imagery is a cognitive function: although there may be concomitant activation of relevant motor or perceptual structures, it mainly serves to drive imagery through the provision of efference copy. Our data point towards a more integrated and ubiquitous role for mental imagery, which does not operate in isolation but in a more dynamic and interactive manner.

Although it is true that I did spend significant time on imagery as a cognitive function, it is not the case that the emulation theory denies the dynamic and interactive role of

imagery that these researchers point out. Indeed, I discussed the role of imagery in selecting motor programs before execution in section 3.2 of the target article – something of which subjects are unaware and which is presumably noncognitive in the relevant sense. (Had these commentators' own take been different, I would have addressed their commentary in the section dealing with data supporting the emulation framework!)

Reed et al. point out that though the cerebellum was implicated in the target article as a likely location for the MSS emulator that is hypothesized to produce motor imagery, deficits in motor imagery can be found in patients with no cerebellar dysfunction, and some cerebellar anomalies don't appear to disrupt motor imagery. The first point is not a problem, as in the emulation theory imagery is not the product of *just* an emulator, but of an emulator being appropriately driven by various motor areas, and so forth. Presumably, patients who have motor imagery deficits have some dysfunction in one of the other components involved in the production of imagery besides the emulator itself. The second point is more difficult, but it hangs on a number of issues, such as the specific regions of the cerebellum damaged and the specific regions involved in emulating the MSS. Again presumably, whatever areas are compromised in the cerebellar atrophy cited by these commentators is not crucially involved in this sort of imagery. But more empirical detail would be needed to say much about this issue.

Tomasino et al. discuss imagery and claim that:

The tasks that are known to tap visual imagery without involving a motor component, such as the Island Test (Kosslyn et al. 1978), the Clock Test (Grossi et al. 1993), the Piazza del Duomo (Bisiach & Luzzatti 1978), have not been considered by the author. Is his omission due to the fact that the classical approach to visual imagery would not easily fit the model?

This is simply inaccurate. Actually, in section 4.5, I did discuss cases of imagery that did not involve any motor action explicitly, and pointed out there that nothing special was involved. Such cases were simply bouts of emulation that lacked a special driving force. A flight simulator still produces instrument readings and visual scenes even if the pilot doesn't do anything with any of the controls but only examines the instruments and display. These kinds of imagery fit the model as easily as any special case fits its generalization.

Tomasino et al. also question the fact that, although the distinction between emulation and simulation was a centerpiece of the discussion of motor imagery, it seemed not to play a role in the discussion of visual imagery:

Earlier in the target article (sect. 2.3), the author states that the simulation theory itself is not sufficient to explain the motor imagery phenomena and claims that an emulator of the musculoskeletal system is needed. When he then turns to discuss the visual imagery domain, it becomes far from clear where the simulation ends and the emulation starts, raising the doubt whether the model is applicable outside the motor domain.

But this too is simply inaccurate. The last paragraph of section 4.3 is focused on exactly this issue, and explains why the distinction between simulation and emulation is much more clear and obvious in the case of visual imagery than in motor imagery. Even when motor commands are involved in visual imagery, as they sometimes are (though not always), the bare motor commands by themselves simply cannot be all there is to the visual imagery. The *same motor command* is used to rotate a visual image clockwise,

whether the image is a capital “L” or a “7” or a “E,” and so the corollary discharge itself *cannot* explain the different images that result. If the motor command was the only thing involved, then, since we have the same motor command, we should have the same image. We don’t have the same image; therefore, the motor command can’t be all there is to it.

Tomasino et al. also take issue with my claim (section 3.2) that the primary motor cortex (M1) “is conspicuously silent during motor imagery,” and they cite evidence to the effect that primary motor cortex may be active during motor imagery. Other than the fact that the commentators misquoted me by leaving out a parenthetical “usually” that I put in between the “is” and the “conspicuously” exactly because of such studies, I am in complete agreement. The involvement of this or that area in any aspect of emulation is an empirical question, and if M1 is involved, then it is involved. The emulation framework is not itself a hypothesis about where its functional components are located.

Gaveau et al. argue:

To make his claim more convincing, Grush has failed to address key issues such as: (1) What, besides the word “emulation,” is common between the predictive activities involved in tasks as different as guiding the hand toward a target (motor control), generating a structured sentence (language), or determining where “Maxi will look” (theory of mind)? (2) What could be the nature of the common substrate that is postulated to be involved in those incredibly dissimilar tasks?

First, note that there need not be a common substrate, since some of these may involve different emulators, in different locations. But let’s look carefully at the examples. In guiding a hand toward a target, visual and kinaesthetic imagery is used to aid the subject’s ongoing perception of the event as it unfolds. While I don’t have any reason to think that generating a structured sentence in itself involves emulation, there is a good deal of evidence that understanding a sentence involves the construction of semantic structures that exploit imagery: As one interprets the sentence, one constructs images (perhaps largely involving amodal imagery) of the scenario being described; and in understanding another’s reasoning, I may construct an imaginative scenario that reconstructs, to some extent, the situation of the other agent and then determine how I would be inclined to react in that situation. Now, in all of these cases, the claim that imagery is involved is substantive and contested, and so I don’t claim to have *established* anything. But it also seems to me entirely plausible to suppose that the *same* imagery processes may be involved. Imagining what my arm looks like when moving might be *exactly* the process involved in (1) perceptual processing, as my arm moves and I watch it; (2) trying to understand what it is like to be someone who is watching his or her arm move; and (3) constructing a meaning for (i.e., understanding) the sentence “I watched my arm move,” as I hear or read it (see Langacker 1999b; Talmy 2000a). Thus, while it has not been established that these tasks employ the same mechanism, it seems to me highly implausible that the only connection between them is a semantic confusion.

Gaveau et al. also remark:

This may suggest an exactly opposite interpretation of the Wexler experiments (Wexler et al. 1998), which are presented as a key support to the emulation theory. How can Grush rule out the possibility that the conflict takes place between the *sensory* outcome predicted by the actual motor command (through the forward model) and the mentally rotated one, and

not between the actual motor command and the command necessary to rotate the object? (emphasis in Gaveau et al.)

I guess I can because it doesn’t seem plausible to suppose that, every time I rotate my wrist, visual images of sideways capital “Ls” are produced (and upside-down backwards “7s,” and all the other images that might conflict with an actual image if I happen to be looking at one). I suspect that I am not understanding the question correctly. Gaveau et al. continue:

It is quite difficult to see how “emulating” the rotation would simply be *possible* when the motor cortex is engaged in a task incompatible with the mental rotation (does this imply the existence of dual forward models?).

One doesn’t need two forward models for this (though I see no reason to rule out the possibility that the brain has more than one forward model for a given task). The interference results from the fact that the motor centers are trying to do two different things: produce one command that actually moves the wrist, and a different command to drive the forward model that produces the imagery (note that we have only one forward model mentioned here). And, as in every case where the motor centers are trying to do two different things, like patting your head and rubbing your belly, there can be degradations in how well either or both of these tasks get executed.

Gaveau et al. go on to cite several lines of evidence that they take to be inconsistent with the emulation framework. The following is typical:

Indeed, dissociations between intact visual imagery and profoundly affected visual perception have been found in several patients (Bartolomeo et al. 1997; Beschin et al. 2000; Goldenberg et al. 1995; Servos et al. 1995). These results openly contradict the notion that visual imagery emerges via an “emulation” of normal vision through top-down processes.

Again, not only do I not see why this should constitute counter-evidence for the emulation theory, it seems clear that the theory predicts exactly this possibility. Perhaps one reason these commentators accuse me of trading on loose analogies and falling into semantic confusions is because they did not carefully look at the diagrams or pay attention to the details of the proposal. For example, a quick look at Figure 6 shows exactly what is happening in such patients. The information flow from vision that would normally be processed in a visual emulator (this flow is represented by lines and boxes between the dotted-lined and dashed-line boxes – the sensory residual, Kalman gain, etc.) has been compromised. But the capacity to drive this emulator via efference copies is intact. The emulation theory *predicts* such dissociations as well as the normal partial sharing of substrates.

R6. Empirical data potentially supporting the emulation framework [Campbell & Pettigrew; Hanakawa, Honda & Hallett (Hanakawa et al.); Hubbard & Ramachandran; Reed, Grubb & Winkielman (Reed et al.); Sadato & Naito; Schubotz & von Cramon]

Campbell & Pettigrew mention some fascinating data that they take to support the emulation framework. Noting that this framework attempts to synthesize the functions of motor control, imagery, and perception, they remark that

“if these three seemingly distinct systems share the same underlying neural mechanisms, then it follows that they must also share a common timing mechanism.” I agree only conditionally with this. If the systems share a neural substrate, then they would share certain temporal features not because there is any sort of master clock, but simply because they are using the same machinery, and indeed are aspects of the same process. Perhaps this is what these commentators mean?

Hubbard & Ramachandran provide a very interesting account of the size-weight illusion that posits an emulator, based in the cerebellum, that models the dynamic characteristics of the objects with which there will be an interaction. This result points to an application that was not covered in the target article – the emulation not just of the MSS *or* of objects in the environment, but of their *interaction*. There are studies that suggest that the cerebellum has models of tools and other entities that are interacted with (see, e.g., Wolpert & Kawato 1998), and presumably the emulator posited to explain the size-weight illusion is of the same sort.

Hanakawa et al. discuss a potential re-interpretation of previously puzzling data based on the idea that certain imagery is amodal and driven by rostral premotor areas (see their commentary for details). Their interpretation of the data seems plausible to me. It also serves as a prod for me to be more specific about the relation between modal and amodal emulation: specifically, whether they will or will not always be driven by the same motor areas serving as the controller. This is, as always, an empirical question, and Hanakawa et al.'s data suggest that there may be a functional division. Their second speculation to the effect that these rostral premotor areas are serving as a controller for an amodal emulator implemented in other structures, strikes me as plausible. The only thing I can add to their fascinating discussion is to point out that the sort of imagery that they classify as amodal, although surely not motor-modality specific, *may* still involve a modal component that is visual – something like a visual presentation of the abacus. But even if this is so, the basic idea that what is happening is the operation of a different emulator – whether it is amodal, or modal in a different modality – is maintained.

The studies discussed by **Reed et al.** – on divers' imagery and emotion recognition – contain results I was not previously aware of. The suggestion that the results on imagery derived from divers of different skill levels bears on the simulation versus emulation theory of imagery issue, is completely clear. Regarding recognition of emotions, I agree that the data are consistent with the emulation account, but, in order to fully make that case, more detail needs to be added to this account. The issues here are entirely similar to those I discuss in Section R10, concerning mirror neurons and imitation.

Sadato & Naito discuss results that suggest that, in fact, the format (or at least one of the formats) of motor imagery is kinaesthetic. The main result is that motor imagery can influence actual (but, importantly, passive) kinaesthesia:

This finding indicates that the emulator, driven by the mental imagery, outputs the “mock” sensory signals in a proprioceptive format, which interferes with the real (but artificially-generated) proprioceptive sensory information from the musculoskeletal system.

This is an ideal paradigm for distinguishing emulation from simulation accounts of imagery, because of the fact that the overt kinaesthetic experience is elicited passively and, hence,

lacks an overt motor component. The only recourse a simulation theorist would have is to say that the motor simulations that are constitutive of motor imagery are in kinematic terms. This is a suggestion made by **Schubotz & von Cramon**, who refer to such terms as a “common neural code.” It is not an implausible position on its own, but it seems to be in conflict with the case of paralyzed phantom limb patients. These patients are of course aware of what kinematic movements their limb ought to be making, even though such movements are not being made. The phenomenology of a phantom limb is not that of an unmovable part for which movement plans cannot be made, such as the earlobe – which does not feel paralyzed (despite the fact that it is not voluntarily independently movable) precisely because no motor plans can be made for earlobes. But this issue is a tricky one, requiring fuller treatment than space considerations allow here. Aside from this slightly controversial position, the **Sadato & Naito** argument is quite convincing, with impressive results, particularly in regard to the articulant-specific results mentioned in the second part of their commentary.

R7. The distinction between modal and amodal emulators [Campbell & Pettigrew; Gärdenfors; Gaveau et al.; Hanakawa et al.; Merfeld; Sathian; Schubotz & von Cramon; Wexler; Wiener & Raab; Wilson]

Merfeld, in the opening paragraph of his commentary, points out that Kalman filtering often takes the amodal form, in part for the same reason as mentioned in the target article (sects. 3.3, 5.1, 6.1): freedom from limitations of the sensory signal format. And **Hanakawa et al.** provide evidence in support of a distinction between modal and amodal imagery. Furthermore, **Campbell & Pettigrew** go on to discuss a fascinating application of the emulation framework, the ability to use a combination of modal and amodal emulators to compensate for an inability to distinguish plant drift from objective environment changes in the calculation of reafference. Finally, **Schubotz & von Cramon** provide references to a number of studies they've conducted, which can be interpreted not only in the generic terms of the emulation framework, but in terms of an interplay between modal and amodal emulators. However, not all commentators were enthusiastic about the distinction.

Gaveau et al. point out that

biasing the input in a given sensory modality leads to an adaptation of that modality (e.g., the waterfall illusion in vision). It is possible that Grush would interpret this result as a change in the emulator (it is a change in the prior probabilities of object motion, which is part of our knowledge of the world – supposedly analogous to the command of a KF). However, in contrast to the prediction of an amodal emulator, it can be shown that this kind of adaptation does not transfer to other modalities.

I'm not sure why this is taken to be evidence against the existence and use of an amodal emulator. If the brain has both modal and amodal emulators, then the unrefined notion of “adaptation” needs to be refined: There will be (a) adaptation that is limited to a modal emulator in a single modality; (b) adaptation that affects the relation between the modal and amodal emulator in a given modality – how they are calibrated, so to speak; and (c) adaptation of the workings of the amodal emulator. For example, after wearing in-

verting prisms, there may be a time at which the modality-specific task of making eye movements in the appropriate direction to foveate a stimulus is correctly relearned, but the stimuli still appear to be on the wrong side. This would involve an adaptation of modality-specific mechanisms while leaving other, amodal mechanisms unadapted. The point is not whether this sort of case actually occurs; rather, the point is that, in the present model, it *can* occur, and so, selective adaptation by itself isn't a problem. In fact, it is predicted by the framework.

Both **Wiener & Raab** and **Gärdenfors** discuss the issue of perception in a way that depends on the issue of the relation between modal and amodal emulation, but it appears that Gärdenfors has understood the proposal much more clearly. While it is true that the emulation model as I articulate it in the target article posits modal emulators that aid in modality specific processing, the remark by Wiener & Raab that "Grush embraces the idea that perception is contingent upon matching the shapes of 'percepts' and of internally generated images (Kosslyn & Sussman 1995)" is not entirely accurate. Perception was clearly described (see the beginning paragraphs of sect. 5 on "Perception") as involving a relation between a modal and amodal emulator, elements of the latter providing interpretations of elements of the former.

Wiener & Raab also discuss what they take to be limitations of modal emulation as exemplified by Mel's Murphy model:

If such two-dimensional arrays of connectionist pixels were to carry geometric projections from spaces of higher dimensionality . . . correct pictorial transformations and object recognition and tracking . . . would be impossible. This consideration renders Grush's image emulator and the corresponding part of the target article's Figure 7 gratuitous.

I agree almost entirely. First it should be pointed out that, while **Wiener & Raab** are making reference to Mel's "Murphy" model (Mel 1988), Mel has another connectionist model (Mel 1986) that does exactly what Wiener and Raab here claim is impossible – it performs accurate predictions of two-dimensional projections of three-dimensional objects during rotation, zoom and pan, even for novel shapes not observed during training. Furthermore, the sort of retinal receptive field remapping demonstrated by Duhamel et al. (Duhamel et al. 1992) appears to demonstrate that such modal emulation is in fact occurring. This remapping is not a theoretical posit but an empirical result from actual single cell recordings. Thus, modal emulation appears to be possible for three-dimensional objects and also to be actually implemented in the brain, its potentially gratuitous nature notwithstanding. Nevertheless, I agree that the modal emulators will have severe limitations, and hence any perception worth the name will necessarily involve interpretation by amodal emulation mechanisms.

Gärdenfors' grasp of the issue is much better. It is exactly the interplay of the modal and amodal emulators that makes perception, which happens in one or more modalities, as contentfully rich as it is, and provides for the ability to correctly anticipate the consequences of movement with three-dimensional objects, and the like.

The point is similar to **Wexler's**, who remarks quite correctly that there are severe limitations to what a visual modality-specific emulator can do (his examples involve sensory aliasing, and are quite similar in content to the cases mentioned by **Wiener & Raab**). He takes the pointing out

of limitations to modality-specific emulation to be a matter of questioning their applicability. However, it seems to me that a more correct conclusion to draw is that they do have application, though this application is limited. What are the possible applications of modal visual emulation, granting its limitations? One, clearly, would be in early visual processing, where the main representations dealt with concern two-dimensional retinal maps and one of the main motor processes is eye movement. For areas of visual processing that are detecting edges and such, the ability to anticipate how these edges and other brightness features will shift as a result of eye movements is not at all trivial. Wexler's argument is:

It would be strange, therefore, if a modal prediction mechanism existed just for those special cases, only to be superseded by a much more general amodal predictor as soon as one steps out of the special case. By parsimony, we might suppose that all predictors are amodal.

But I don't think it would be strange at all. In fact, given the complexity, and long, quirky evolutionary history of the CNS, and the many kinds of sensory systems it has had and dropped or modified in its evolution, it would be surprising if there *were* just one strategy used for all visual processing. The parsimony argument has little pull with me. **Wexler** himself points to the correct cases where modal emulation, despite its limitations, has application and has been implemented. The Duhamel et al. (1992) result that both Wexler and I mention seems to be verification of exactly this sort of modality specific mechanism.

Similar remarks are appropriate to **Schubotz & von Cramon's** commentary. They point out that an "introspectively compelling reason for suggesting independent modal emulation is that some kinds of modal imagery (e.g., a vase) feel purely visual and not at all motor." They then go on to argue, quite correctly, that introspection is a poor guide. However, we don't need to fall back on introspection to motivate the existence of a visual modal emulator. Saccade control and early visual anticipations can not only exploit such modal imagery, but the Duhamel et al. result appears to verify its existence and determine one location. I fully agree with Schubotz & Cramon's point that modal and amodal emulation may not be independent; but I don't think it follows that they are conceptually dependent. I can imagine simple nervous systems with visual systems that exploit emulation limited to the visual modality, and, perhaps because they have no depth perception, do not use such predictions for anything multimodal or amodal, but only for helping to control eye movements and anticipating (imperfectly) the results of such movements as an aid to processing.

Wilson also worries about the distinction between modal and amodal imagery with reference to what she calls the "Images-aren't-pictures problem," and states that:

Grush takes at face value the idea that modality-specific imagery is equivalent to unprocessed sensory input. In fact, this idea has been vigorously challenged.

Although it is true that in my account modality-specific imagery is equivalent to raw sensory input (it is in the same format), it is an open question whether anyone ever can generate *that* kind of imagery by itself. It could very well be that what we think of as "visual imagery" is always a *composite* of both modal and amodal imagery. And the modal part might involve more than simply a measurement of an

amodal component, but might also be supplemented with remembered images. I don't think that Wilson and I are actually disagreeing about anything substantive.

Sathian recommends the following:

Rather than "amodal" emulators, I suggest invoking "multisensory" emulators to provide the link between modality-specific systems and between these systems and abstract representations. I must emphasize that this is not a merely semantic distinction. By "multisensory," I mean a system that receives inputs from more than one sensory modality.

Sathian suggests that the notion of amodal representation is best reserved for linguistic and conceptual systems, and not for anything that could be involved in an emulatory process. I have two replies to this. First, by the definition given in the quote above, what I called "amodal" emulators are (potentially, at least) multisensory. The amodal emulator is posited to work in concert with one or more modal emulators and to get input in the form of residual corrections from one or more modalities. It is true that I provided explicit examples of only an amodal emulator working with one modality, but in the penultimate paragraph of section 4.4 the possibility of more than one modality was discussed, and in fact a number of the same functions were attributed to the amodal emulator as Sathian attributes to a multisensory system. Indeed, the sort of processing taking place in the posterior parietal cortex and studied by Andersen and colleagues (e.g., Cohen & Andersen 2002; Xing & Andersen 2000) is constitutive of what I take to be the function of the central amodal emulator (as an emulator that represents an organism's behavioral space). One reason for calling such a system amodal rather than multimodal is that there are cases where an object cannot currently be sensed by any sensory modality (because it is behind an occluder, is silent and odorless, etc.), yet it is represented as being at a location. I think it is safe to say that our representation of our own behavioral (egocentric) space allows for this, and it is not clear how a multisensory system, in which tags for specific modalities were always present, could accomplish this.

My second point in response to **Sathian's** remarks is that I take "conceptual" and "linguistic" representations to be much more closely tied to sensorimotor behavior than is often recognized, and so there is a theoretical reason to keep amodal representations closely tied to behavior. For a more detailed discussion on this, I point the interested reader towards the work of Ronald Langacker and Leonard Talmy (cf. especially Langacker 1999b; Talmy 2000a).

R8. Where in the CNS are emulators implemented? [Campbell & Pettigrew; Donchin & Raz; Hanakawa et al.; Hubbard & Ramachandran; Reed et al.]

First, to follow up a point made in section R2: the emulation framework does not posit *one single* emulator, and so asking for the neural substrate of *the* emulator is not anything I can do. For example, **Donchin & Raz's** question about neural substrates is prefaced by the following remark: "Grush implies that this general network manifests in converging neurophysiological mechanisms." This may be true for some cases, but it may be false for others. Visual imagery and visual perception probably share a number of neural bases, but other applications of emulation may not.

With regard to the musculoskeletal emulator that is

posited to subserve certain motor control functions and motor imagery, both **Donchin & Raz** and **Campbell & Pettigrew** mention the cerebellum and associated hindbrain structures. As I mentioned in the target article, these are areas which have been converged upon by some earlier neurophysiologists to specifically posit forward models (e.g., Ito 1970; 1984; Kawato 1989; 1990).

I am grateful to **Donchin & Raz** for the additional references implicating the cerebellum in various functions for which the emulation theory seems to have potential application. The commentary by **Hubbard & Ramachandran** also implicates the cerebellum, again, not just for the MSS emulator, but apparently for an emulator that models the dynamic characteristics of objects in the environment and how they might interact with the body.

It remains, however, that I am not primarily interested in making guesses about neural localization, at least not at this stage – though this may irritate some people, including many commentators here. The brain is extremely complicated, and with a few exceptions, trying to nail down areas of neural implementation is difficult. I take my task to be one of helping to articulate, in a clear way, some of the information processing structures that can, in turn, be an aid to others who are interested in, and have competence and experience with, ferreting out neural localizations. The **Hanakawa et al.** commentary is a perfect example. My goal was to introduce the emulation framework in such a way that those who are in a position to responsibly do the neuroscience will be better able to assess its applicability and discern its implementation. Science is a team effort, and not every member of the team has to be competent in all the jobs that other members of the team perform.

R9. What level of detail should emulators represent? [Gärdenfors; Stein; Stojanov & Bickhard; Webb]

The issue of the level of detail present or required in emulations (something I mentioned quite briefly in sect. 5.2 of the target article) came up in a number of commentaries. For example, **Stein** describes (in a way that corrected some misunderstandings on my part) MetaToto's emulating processes as being schematic, this schematicity being a function of the relation between imagination and its memory, which represented only salient features of the environment. I am in complete agreement with Stein's suggestion that it is most likely the salient features of the target system (where "salient" means "salient for the sensorimotor exigencies of the organism") that are typically emulated.

With the recognition of schematic emulation in hand, we can address an objection by **Stojanov & Bickhard**. These authors bring up change blindness, in which many subjects can fail to notice what are in fact significant changes. As Stojanov & Bickhard put it:

According to Grush's KF framework, because the estimate does not match the stimuli, the Kalman gain should increase, which would lead to an accurate representation and perception of the changed photo.

In these cases, it seems that one of two things is happening. Either the emulation is schematic and is not representing a specific (or specific enough) value for the changed element, and so there is in fact not a mismatch. The other possibility

is that attentional mechanisms of some sort are crucial in whatever process is playing the role of Kalman gain, such that even if some element is being represented in detail at some level of processing, contradictory sensory input may be ignored in some cases of residual correction. That is, residual corrections in the actual case may not be a matter of strictly defined Kalman gain, but a function of other factors, including attention.

However, **Webb's** question is pertinent here: "how 'schematic' . . . can the emulator be, yet still count as an emulator?" It is one thing to say that emulation can be schematic, quite another to draw a line between increasingly schematic emulations and any kind of anticipation. I'd really like to have some good, theoretically motivated way to make a distinction between information processing structures that are genuinely emulation-involving and those that are not (not because I think it is important to distinguish between "higher" and "lower" organisms, but because I would like to distinguish representational from non-representational processes – I have no stake in the claim that simple nervous systems can't represent). What is important is to keep our theories clear, and do the empirical work. Perhaps **Webb** is right when she suggests that "we are led to consider a continuum, from complete and detailed internal models of the plant, environment, and the measurement process, to extremely simple processes that might suffice for prediction." But, on the other hand, there may be some discontinuities, as **Gärdenfors** seems to assume when he says "An important question for future research then becomes: Why do humans have all these, apparently very successful, emulators for causes and a theory of mind, and why do other species not have them?" What we all seem to agree on is that these are empirical questions.

As I mentioned, a distinction needs to be made here: I think that **Stein** and **Webb** are, initial appearances notwithstanding, focusing on two different and orthogonal dimensions of schematicity. In the dimension picked out by **Stein**, there is no question but that the features of some target system (the environment in this case) are being emulated. What is more or less schematic is the extent of detail that the target system is represented as having. The dimension **Webb** is focusing on has to do with the extent to which the product of the process that is driven by a corollary discharge can be adequately interpreted as being a representation of a target system at all, as opposed to some signal that is useful in preparing the organism for an impending event (e.g., in the way that preparation for motor activity can engage metabolic changes, such as increased heart rate, even before the activity begins, but these changes don't represent or emulate the activity). I suppose that these two continua are both present, and that perhaps there is some correlation between them. But, beyond recognizing the interest of the questions posed, I have nothing to add at this point. I wish I did.

R10. Emulation, imitation, and mirror neurons [Campbell & Pettigrew; Reed et al.; Slaughter; Wexler; Wilson]

Finally, several commentators (**Campbell & Pettigrew**; **Slaughter**; **Wilson**) have discussed the relation between the emulation framework (especially emulators of one's own body) and phenomena such as mirror neurons, imita-

tion, ideomotor apraxia, and so on. I agree with these commentators that this is an area of potentially great interest, and if it turns out that the emulation theory can help shed light on these phenomena, then I will be quite pleased. I made an attempt at such connections in **Grush 1995**, Chapter 5, but have subsequently become dissatisfied with some of the details of that treatment (though, to my knowledge, this is the first even remotely detailed attempt to articulate connections between forward models as used in motor control and "other minds" phenomena). I share the intuition that there must be some connection, but I'm unsure of how the details would work in information processing terms, and getting the details of the information processing structure down is my main concern (much to the disappointment of those interested in locating neural substrates). It has recently come to my attention that **Susan Hurley** has been working on the connections between emulation theory, imitation, and mirror neurons (**Hurley forthcoming a**; forthcoming b).

I found **Wilson's** suggestion to the effect that one purpose of such processing might be to aid the perception of biomechanical motion to be extremely interesting. As described, the emulation framework would claim that the increased capacity to accurately perceive biomechanical movement would accrue by having a process model of biomechanical motion to serve in some sort of KF-like filtering process. But it is not clear that a model of *one's own* biomechanical movement – presented, as it were, from the inside – can be used to shed light on externally observed biomechanical motion – presented from the third person perspective. Perhaps one model can serve both, I'm not sure. (This worry is similar to **Wexler's** worry about the relation between the MSS emulator and the environment emulator. And it is akin to **Reed et al.'s** remarks concerning emotion perception.)

Another possibility, sparked by **Wilson's** suggestion but still extremely tentative and poorly specified, is as follows. Perhaps there are two different but related biomechanical emulators: one tied to one's own body in a first-person sort of way (the MSS emulator as described in the target article), and a second dedicated to representing biomechanical movement *as observed from the outside*. These would presumably be related because one thing a higher organism does when it moves about is to keep track of how its body and movements will appear to others; and so the same efference copies that are used to drive an emulator of my internal biomechanics could also drive an emulator of my "external" biomechanics in order to update how, in my estimation, my body and its environmental compartment looks to others. But since the external biomechanical emulator need not be exclusively tied to the task of maintaining a representation of how *I* will look to others, it is capable of being used to model the biomechanics of other individuals for perceptual processing. And then, because of the connections between the external biomechanical emulator and motor areas that is present for the first function, the connection becomes manifest in the second function as well. This is a particularly speculative idea. However, this topic will undoubtedly be the focus of intensive research over the coming years, so it is perhaps best to end the speculations here.

References

Letters “a” and “r” appearing before authors’ initials refer to target article and response, respectively.

- Adamovich, S. V., Archambault, P. S., Ghafouri, M., Levin, M. F., Poizner, H. & Feldman, A. G. (2001) Hand trajectory invariance in reaching movements involving the trunk. *Experimental Brain Research* 138:288–303. [RB]
- Adler, B., Collewijn, H., Curio, G., Grusser, O. J., Pause, M., Schreier, U. & Weiss, L. (1981) Sigma-movement and sigma-nystagmus: A new tool to investigate the gaze-pursuit system and visual-movement perception in man and monkey. *Annals of the New York Academy of Sciences* 374:284–302. [TGC]
- Adolphs, R., Damasio, H., Tranel, D., Cooper, G. & Damasio, A. R. (2000) A role for somatosensory cortices in the visual recognition of emotion as revealed by three-dimensional lesion mapping. *Journal of Neuroscience* 20:2683–90. [CLR]
- Alain, C., Arnott, S. R., Hevenor, S., Graham, S. & Grady, C. L. (2001) “What” and “where” in the human auditory system. *Proceedings of the National Academy of Science USA* 98(21):12301–306. [aRG]
- Alexander, G. E. & Crutcher, M. D. (1990) Neural representations of the target (goal) of visually guided arm movements in three motor areas of the monkey. *Journal of Neurophysiology* 64:164–78. [BT]
- Ariff, G., Donchin, O., Nanayakkara, T. & Shadmehr, R. (2002) A real-time state predictor in motor control: Study of saccadic eye movements during unseen reaching movements. *Journal of Neuroscience* 22:7721–29. [OD]
- Ashe, J., Taira, M., Smyrnis, N., Pellizzer, G., Georgakopoulos, T., Lurito, J. T. & Georgopoulos, A. P. (1993) Motor cortical activity preceding a memorized movement trajectory with an orthogonal bend. *Experimental Brain Research* 95:118–30. [BT]
- Balasubramaniam, R. & Feldman, A. G. (2004) Guiding movements without redundancy problems. In: *Coordination dynamics: Issues and trends*, ed. V. K. Jirsa & J. A. S. Kelso. Springer. [RB]
- Barsalou, L. W. (1999) Perceptual symbol systems. *Behavioral and Brain Sciences* 22(4):577–609. [aRG, CLR]
- Barsalou, L. W., Simmons, W. K., Barbey, A. K. & Wilson, C. D. (2003) Grounding conceptual knowledge in modality-specific systems. *Trends in Cognitive Science* 7(2):84–91. [VG]
- Barsalou, L., Solomon, K. O. & Wu, L. (1999) Perceptual simulation in conceptual tasks. In: *Cultural, typological, and psychological perspectives in cognitive linguistics*, ed. M. K. Hiraga, C. Sinha & S. Wilcox. John Benjamins. [aRG]
- Bartolomeo, P., Bachoud-Levi, A. C. & Denes, G. (1997) Preserved imagery for colours in a patient with cerebral achromatopsia. *Cortex* 33:369–78. [VG, rRG]
- Batista, A. P., Buneo, C. A., Snyder, L. H. & Andersen, R. A. (1999) Reach plans in eye-centered coordinates. *Science* 285:257–60. [OD]
- Behrmann, M. (2000) The mind’s eye mapped onto the brain’s matter. *Trends in Psychological Science* 9(2):50–54. [aRG]
- Bell, C., Bodznick, D., Montgomery, J. & Bastian, J. (1997) The generation and subtraction of sensory expectations within cerebellum-like structures. *Brain, Behavior and Evolution* 50 (Suppl. 1):17–31. [TGC]
- Bell, C. C., Libouban, S. & Szabo, T. (1983) Pathways of the electric organ discharge command and its corollary discharges in mormyrid fish. *Journal of Comparative Neurology* 216(3): 327–38. [TGC]
- Bellman, R. (1964) *Perturbation techniques in mathematics, physics, and engineering*. Holt, Rinehart and Winston. [VGo]
- Benson, D. F. (1994) *The neurology of thinking*. Oxford Press. [RIS]
- Bernstein, N. A. (1947) *On the construction of movements*. Medgiz. [MLL]
- (1967) *The coordination and regulation of movements*. Pergamon. [MLL, RB]
- Berthoz, A. (1996) The role of inhibition in the hierarchical gating of executed and imagined movements. *Brain Research. Cognitive Brain Research* 3:101–13. [OD]
- Beschin, N., Basso, A. & Della Sala, S. (2000) Perceiving left and imagining right: Dissociation in neglect. *Cortex* 36:401–14. [VG, rRG]
- Bickhard, M. H. (1980) *Cognition, convention, and communication*. Praeger. [GS]
- (1993) Representational content in humans and machines. *Journal of Experimental and Theoretical Artificial Intelligence* 5:285–333. [GS]
- (2000) Motivation and emotion: An interactive process model. In: *The caldron of consciousness*, ed. R. D. Ellis & N. Newton, pp. 161–78. John Benjamins. [GS]
- (2004) Process and emergence: Normative function and representation. *Axiomathes* 14:135–69. [GS]
- Bickhard, M. H. & Campbell, R. L. (1996) Topologies of learning and development. *New Ideas in Psychology* 14(2):111–56. [GS]
- Bickhard, M. H. & Terveen, L. (1995) *Foundational issues in artificial intelligence and cognitive science: Impasse and solution*. Elsevier. [GS]
- Bisiach, E. & Luzzatti, C. (1978) Unilateral neglect of representational space. *Cortex* 14:129–33. [rRG, BT]
- Blakemore, S. J., Frith, C. D. & Wolpert, D. M. (2001) The cerebellum is involved in predicting the sensory consequences of action. *NeuroReport* 12(9):1879–84. [TGC]
- Blakemore, S. J., Goodbody, S. J. & Wolpert, D. M. (1998) Predicting the consequences of our own actions: The role of sensorimotor context estimation. *The Journal of Neuroscience* 18(18):7511–18. [aRG]
- Blakemore, S. J., Wolpert, D. M. & Frith, C. D. (1998) Central cancellation of self-produced tickle sensation. *Nature Neuroscience* 1(7):635–40. [TGC]
- (2002) Abnormalities in the awareness of action. *Trends in Cognitive Sciences* 6(6):237–42. [TGC]
- Bodznick, D., Montgomery, J. C. & Carey, M. (1999) Adaptive mechanisms in the elasmobranch hindbrain. *Journal of Experimental Biology* 202:1357–64. [TGC]
- Borah, J., Young, L. R. & Curry, R. E. (1988) Optimal estimator model for human spatial orientation. *Annals of the New York Academy of Sciences* 545:51–73. [DMM]
- Boring, E. G. (1950) *History of experimental psychology*. Appleton-Century-Crofts. [JSJ]
- Botwinick, M. & Cohen, J. (1998) Rubber hands “feel” touch that eyes see. *Nature* 391:756. [BT]
- Brooks, R. A. (1986) A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation* 2:14–23. [aRG]
- (1991) Intelligence without representation. *Artificial Intelligence* 47:139–60. [aRG]
- Brugger, P., Kollias, S. S., Muri, R. M., Crelier, G., Hepp-Reymond, M. C. & Regard, M. (2000) Beyond re-membering: Phantom sensations of congenitally absent limbs. *Proceedings of the National Academy of Science USA* 97:6167–72. [BT]
- Brunet, E., Sarfati, Y., Hardy-Bayle, M. C. & Decety, J. (2000) A PET investigation of the attribution of intentions with a nonverbal task. *Neuroimage* 11:157–66. [OD]
- Bryson, A. & Ho, Y.-C. (1969) *Applied optimal control; Optimization, estimation, and control*. Blaisdell. [aRG]
- Buneo, C. A., Jarvis, M. R., Batista, A. P. & Andersen, R. A. (2002) Direct visuomotor transformations for reaching. *Nature* 416:632–36. [KS]
- Calvert, G. A. (2001) Cross-modal processing in the human brain: Insights from functional neuroimaging studies. *Cerebral Cortex* 11:1111–23. [KS]
- Calvo Garzón, F. (Submitted) Towards a general theory of antirepresentationalism. [FCG]
- Campbell, T. G., Erickson, G., Wallis, G., Liu, G. B. & Pettigrew, J. D. (2003) Correlated individual variation of efference copy and perceptual rivalry timing. Program No. 550.1. 2003 *Abstract Viewer/Itinerary Planner*. Society for Neuroscience. Online Publication. [TGC]
- Carpenter, A. F., Georgopoulos, A. P. & Pellizzer, G. (1999) Motor cortical encoding of serial order in a context-recall task. *Science* 283:1752–57. [BT]
- Castelli, F., Frith, C., Happe, F. & Frith, U. (2002) Autism, Asperger syndrome and brain mechanisms for the attribution of mental states to animated shapes. *Brain* 125:1839–49. [OD]
- Catalan, M. J., Honda, M., Weeks, R. A., Cohen, L. G. & Hallett, M. (1998) The functional neuroanatomy of simple and complex sequential finger movements: A PET study. *Brain* 121(2):253–64. [BT]
- Charpentier, A. (1891) Analyse experimentale de quelques elements de la sensation de poids. *Archives de Physiologie Normales et Pathologiques* 3:122–35. [EMH]
- Chen, R., Cohen, L. G. & Hallett, M. (1997) Role of the ipsilateral motor cortex in voluntary movement. *Canadian Journal of Neurological Science* 24:284–91. [BT]
- Chen, W., Kato, T., Zhu, X. H., Ogawa, S., Tank, D. W. & Ugurbil, K. (1998) Human primary visual cortex and lateral geniculate nucleus activation during visual imagery. *NeuroReport* 9(16):3669–74.
- Christensen, W. D. & Bickhard, M. H. (2002) The process dynamics of normative function. *Monist* 85(1):3–28. [GS]
- Clark, A. (1997) *Being there: Putting brain, body and world together again*. MIT Press. [JSJ, GS]
- (2003) *Natural-born cyborgs*. Oxford University Press. [TD]
- Clark, A. & Chalmers, D. (1998) The extended mind. *Analysis* 58(1):7–19. [TD, rRG]
- Cohen, M. S., Kosslyn, S. M., Breiter, H. C., Digirolamo, G. J., Thompson, W. L., Anderson, A. K., Bookheimer, S. Y., Rosen, B. R. & Belliveau, J. W. (1996) Changes in cortical activity during mental rotation: A mapping study using functional MRI. *Brain* 119:89–100. [BT]
- Cohen, Y. E. & Andersen, R. A. (2002) A common reference frame for movement plans in the posterior parietal cortex. *Nature Review Neuroscience* 3:553–62. [rRG, KS]
- Cooper, L. A. & Shepard, R. N. (1973) Chronometric studies of the rotation of mental images. In: *Visual information processing*, ed. W. G. Chase, pp. 24–58. Academic Press. [TD]
- Courchesne, E. (1997) Brainstem, cerebellar and limbic neuroanatomical abnormalities in autism. *Current Opinion in Neurobiology* 7:269–78. [OD]

- Craik, K. (1943) *The nature of explanation*. Cambridge University Press. [aRG]
- Craske, B. (1977) Perception of impossible limb positions induced by tendon vibration. *Science* 196:71–73. [MLL, NS]
- Damasio, A. R. (1989) Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition* 33:25–62. [CLR]
- (1994) *Descartes' error: Emotion, reason, and the human brain*. Putnam. [FCC, aRG]
- Danckert, J., Ferber, S., Doherty, T., Steinmetz, H., Nicolle, D. & Goodale, M. A. (2002) Selective, non-lateralized impairment of motor imagery following right parietal damage. *Neurocase* 8:194–204. [VG]
- Dartnall, T. H. (2003) Externalism extended. In: *Proceedings of the Joint Fourth International Conference on Cognitive Science and the Seventh Australasian Society for Cognitive Science Conference, University of New South Wales, Sydney, Australia, July 2003*, ed. P. Szlezak, pp. 94–99. The University of New South Wales Press. [TD]
- Decety, J. & Jeannerod, M. (1995) Mentally simulated movements in virtual reality: Does Fitts' law hold in motor imagery? *Behavioral Brain Research* 72:127–34. [OD, aRG]
- Decety, J., Perani, D., Jeannerod, M., Bettinardi, V., Tadary, B., Woods, R., Mazziotta, J. C. & Fazio, F. (1994) Mapping motor representations with positron emission tomography. *Nature* 371:600–602. [BT]
- Deiber, M.-P., Ibanez, V., Honda, M., Sadato, N., Raman, R. & Hallett, M. (1998) Cerebral processes related to visuomotor imagery and generation of simple finger movements studied with positron emission tomography. *Neuroimage* 7(2):73–85. [aRG, TH]
- de Renzi, E., Motti, F. & Nichelli, P. (1980) Imitating gestures – A quantitative approach to the ideomotor apraxia. *Archives of Neurology* 37:6–10. [BT]
- Desmurget, M. & Grafton, S. (2000) Forward modeling allows feedback control for fast reaching movements. *Trends in Cognitive Sciences* 4(11):423–31. [aRG]
- Dijkstra, T. M., Schonerg, G. & Gielen, C. C. (1994) Temporal stability of the action-perception cycle for postural control in a moving visual environment. *Experimental Brain Research* 97:477–86. [MLL]
- di Pellegrino, G., Ladavas, E. & Farnè, A. (1997) Seeing where your hands are. *Nature* 21(388): 730. [BT]
- Donald, M. (1994) Precis of the origins of modern mind: Three stages in the evolution of culture and cognition. *Behavioral and Brain Sciences* 16:737–91. [HW]
- Donchin, O., Francis, J. T. & Shadmehr, R. (2003) Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: Theory and experiments in human motor control. *Journal of Neuroscience* 23:9032–45. [OD]
- Dretske, F. I. (1988) *Explaining behavior*. MIT Press. [GS]
- Droulez, J. & Comilleau-Peres, V. (1993) Application of the coherence scheme to the multisensory fusion problem. In: *Multisensory control of movement*, ed. A. Berthoz, pp. 485–501. Oxford University Press. [DMM]
- Droulez, J. & Darlot, C. (1989) The geometric and dynamic implications of the coherence constraints in three-dimensional sensorimotor interactions. In: *Attention and performance, vol. XIII*, ed. M. Jeannerod, pp. 495–526. Erlbaum. [DMM]
- Duhamel, J.-R., Colby, C. & Goldberg, M. E. (1992) The updating of the representation of visual space in parietal cortex by intended eye movements. *Science* 255(5040):90–92. [PG, arRG, MW]
- Eccles, J. C. (1979) Introductory remarks. In: *Cerebro-cerebellar interactions*, ed. J. Massion & K. Sasaki, pp. 10–18. Elsevier. [JSJ]
- Eliasmith, C. & Anderson, C. (2003) *Neural engineering: Computational, representation, and dynamics in neurobiological systems*. MIT Press. [aRG]
- Ellis, R. D. (1995) *Questioning consciousness*. Benjamins. [NN]
- Enoka, R. M. (1994) *Neuromechanical basis of kinesiology*. Human Kinetics. [MLL]
- Erlhagen, W. & Schöner, G. (2002) Dynamic field theory of movement preparation. *Psychological Review* 109:545–72. [FCCG]
- Fadiga, L., Fogassi, L., Gallese, V. & Rizzolatti, G. (2000) Visuomotor neurons: Ambiguity of the discharge or “motor” perception? *International Journal of Psychophysiology* 35:165–77. [RIS]
- Falchier, A., Clavagnier, S., Barone, P. & Kennedy, H. (2002) Anatomical evidence of multimodal integration in primate striate cortex. *Journal of Neuroscience* 22:5749–59. [KS]
- Farah, M. J., Hammond, K. M., Levine, D. N. & Calvanio, R. (1988) Visual and spatial mental imagery: Dissociable systems of representation. *Cognitive Psychology* 20(4):439–62. [aRG]
- Farah, M. J., Soso, M. J., Dasheiff, R. M. (1992) Visual angle of the mind's eye before and after unilateral occipital lobectomy. *Journal of Experimental Psychology: Human Perception and Performance* 18(1):241–46. [aRG]
- Fauconnier, G. (1985) *Mental spaces: Aspects of meaning construction in natural language*. MIT Press. [aRG]
- Feinberg, I. (1978) Efference copy and corollary discharge: Implications for thinking and its disorders. *Schizophrenia Bulletin* 4(4):636–40. [TGC]
- Feldman, A. G. (1986) Once more on the equilibrium-point hypothesis (λ model) for motor control. *Journal of Motor Behavior* 18:17–54. [MLL]
- Feldman, A. G. & Latash, M. L. (1982) Afferent and efferent components of joint position sense: Interpretation of kinaesthetic illusions. *Biological Cybernetics* 42:205–14. [MLL]
- Feldman, A. G. & Levin, M. F. (1993) Control variables and related concepts in motor control. *Concepts in Neuroscience* 4:25–51. [rRG]
- (1995) The origin and use of positional frames of reference in motor control. *Behavioral and Brain Sciences* 18(4):723–806. [RB, rRG, MLL]
- Feltz, D. L. & Landers, D. M. (1983) The effects of mental practice on motor skill learning and performance: A meta-analysis. *Journal of Sport Psychology* 5:25–57. [NS]
- Feynman, R. P. (2001) *The pleasure of finding things out, 1st edition*. Penguin Books. [TGC]
- Flanagan, J. R. & Beltzner, M. A. (2000) Independence of perceptual and sensorimotor predictions in the size-weight illusion. *Nature Neuroscience* 3:737–41. [EMH]
- Flavell, J. H. (1999) Cognitive development: Children's knowledge about the mind. *Annual Review of Psychology* 50:21–45. [aRG]
- Freides, D. (1974) Human information processing and sensory modality: Cross-modal functions, information complexity and deficit. *Psychological Bulletin* 81:284–310. [KS]
- Freyd, J. J. & Finke, R. A. (1984) Representational momentum. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 10:126–32. [MWi]
- Freyd, J. J. & Jones, K. T. (1994) Representational momentum for a spiral path. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20:968–76. [MWi]
- Frith, C. D., Blakemore, S. & Wolpert, D. M. (2000) Explaining the symptoms of schizophrenia: Abnormalities in the awareness of action. *Brain Research. Brain Research Reviews* 31:357–63. [OD]
- Frith, C. D. & Gallagher, S. (2002) Models of the pathological mind. *Journal of Consciousness Studies* 9:57–80. [OD]
- Frith, U. (2001) Mind blindness and the brain in autism. *Neuron* 32:969–79. [OD]
- Gallagher, S. & Meltzoff, A. (1996) The earliest sense of self and others: Merleau-Ponty and recent developmental studies. *Philosophical Psychology* 9:213–36. [VS]
- Ganis, G., Keenan, J. P., Kosslyn, S. M. & Pascual-Leone, A. (2000) Transcranial magnetic stimulation of primary motor cortex affects mental rotation. *Cerebral Cortex* 10:175–80. [BT]
- Cärdenfors, P. (2003) *How homo became sapiens: On the evolution of thinking*. Oxford University Press. [PG]
- Celb, A. (1974) *Applied optimal estimation*. MIT Press. [aRG]
- Georgopoulos, A. P., Lurito, J. T., Petrides, M., Schwartz, A. B. & Massey, J. T. (1989) Mental rotation of the neuronal population vector. *Science* 243:234–36. [BT]
- Gerardin, E., Sirigu, A., Lehericy, S., Poline, J. B., Gaymard, B., Marsault, C., Agid, Y. & Jeannerod, M. (2000) Partially overlapping neural networks for real and imagined hand movements. *Cerebral Cortex* 10:1093–104. [BT]
- Gerloff, C., Corwell, B., Chen, R., Hallett, M. & Cohen, L. G. (1998) The role of the human motor cortex in the control of complex and simple finger movement sequences. *Brain* 121:1695–709. [BT]
- Ceyer, S., Matelli, M., Luppino, G. & Zilles, K. (2000) Functional neuroanatomy of the primate isocortical motor system. *Anatomy and Embryology* 202(6):443–74. [TH]
- Gibson, J. J. (1966) *The senses considered as perceptual systems*. Houghton Mifflin. [EC]
- (1979/1986) *The ecological approach to visual perception*. Houghton Mifflin/Erlbaum. [JSJ, RIS]
- Glasauer, S. (1992) Interaction of semicircular canals and otoliths in the processing structure of the subjective zenith. *Annals of the New York Academy of Sciences* 656:847–49. [DMM]
- Glezer, V. D., Gauzelman, V. E. & Shcherbach, T. A. (1985) Relationship between spatial and spatial-frequency characteristics of receptive fields of cat visual cortex. *Neuroscience and Behavioral Physiology* 15(6):511–19. [VGo]
- Goldenberg, G., Mullbacher, W. & Nowak, A. (1995) Imagery without perception – a case study of anosognosia for cortical blindness. *Neuropsychologia* 33:1373–82. [VG, rRG]
- Goodnow, J. J. & Levine, R. A. (1973) “The grammar of action”: Sequence and syntax in children's copying. *Cognitive Psychology* 4:82–98. [ADS]
- Goodwin, G. M., McCloskey, D. I. & Matthews, P. B. C. (1972a) Proprioceptive illusions induced by muscle vibration: Contribution by muscle spindles to perception? *Science* 175:1382–84. [NS]
- (1972b) The contribution of muscle afferents to kinesthesia shown by vibration induced illusions of movement and by the effects of paralyzing joint afferents. *Brain* 95:705–48. [NS]

- Copnik, A. (1998) Explanation as orgasm. *Minds and Machines* 8:101–18. [PG]
- Gordon, R. M. (1986) Folk psychology as simulation. *Mind and Language* 1:158–71. [aRG]
- Grafton, S. T., Arbib, M. A., Fadiga, L. & Rizzolatti, G. (1996) Localization of grasp representations in humans by positron emission tomography. *Experimental Brain Research* 112:103–11. [BT]
- Grafton, S. T., Hazeltine, E. & Ivry, R. (1995) Functional mapping of sequence learning in normal humans. *Journal of Cognitive Neuroscience* 7:497–510. [BT]
- (1998) Abstract and effector-specific representations of motor sequences identified with PET. *Journal of Neuroscience* 18:9420–28. [BT]
- Graziano, M. S. A. (1999) Where is my arm? Relative role of vision and proprioception in the neural representation of limb position. *Proceedings of the National Academy of Sciences USA* 96:10418–21. [BT]
- Grea, H., Pisella, L., Rossetti, Y., Desmurget, M., Tilkete, C., Grafton, S., Prablanc, C. & Vighetto, A. (2002) A lesion of the posterior parietal cortex disrupts on-line adjustments during aiming movements. *Neuropsychologia* 40(13):2471–80. [VG]
- Grossi, D., Angelini, R., Pecchinenda, A. & Pizzamiglio, L. (1993) Left imaginal neglect in hemianattention: Experimental study with the o'clock test. *Behavioural Neurology* 6:155–58. [rRG, BT]
- Grush, R. (1995) *Emulation and cognition*. Doctoral Dissertation, Department of Cognitive Science and Philosophy, University of California, San Diego. UMI. [arRG, EMH]
- (1997) The architecture of representation. *Philosophical Psychology* 10(1):5–25. [rRG]
- (1998) Wahrnehmung, Vorstellung und die sensorimotorische Schleife. (English translation: Perception, imagery, and the sensorimotor loop) In: *Bewußtsein und Repräsentation*, ed. F. Esken & H.-D. Heckmann. Verlag Ferdinand Schöningh. [PG]
- (2000) Self, world and space: The meaning and mechanisms of ego- and allocentric spatial representation. *Brain and Mind* 1(1):59–92. [rRG]
- (2001) The semantic challenge to computational neuroscience. In: *Theory and method in the neurosciences*, ed. P. Machamer, R. Grush & P. McLaughlin. University of Pittsburgh Press. [rRG]
- (2003) In defense of some “Cartesian” assumptions concerning the brain and its operation. *Biology and Philosophy* 18(1):53–93. [rRG]
- Grusser, O. J. (1995) On the history of the ideas of efference copy and reafference. *Clio Medica* 33:35–55. [TGC]
- Guérin, F., Ska, B. & Belleville, S. (1999) Cognitive processing of drawing abilities. *Brain and Cognition* 40:464–78. [ADS]
- Haarmeier, T., Thier, P., Repnow, M. & Petersen, D. (1997) False perception of motion in a patient who cannot compensate for eye movements. *Nature* 389(6653):849–52. [TGC]
- Haken, H., Kelso, J. A. S. & Bunz, H. (1985) A theoretical model of phase transitions in human hand movements. *Biological Cybernetics* 51:347–56. [CBW]
- Hanakawa, T., Honda, M., Okada, T., Fukuyama, H. & Shibasaki, H. (2003a) Differential activity in the premotor cortex subdivisions in humans during mental calculation and verbal rehearsal tasks: A functional magnetic resonance imaging study. *Neuroscience Letters* 347(3):199–201. [TH]
- (2003b) Neural correlates underlying mental calculation in abacus experts: A functional magnetic resonance imaging study. *Neuroimage* 19(2, Pt. 1):296–307. [TH]
- Hanakawa, T., Honda, M., Sawamoto, N., Okada, T., Yonekura, Y., Fukuyama, H. & Shibasaki, H. (2002) The role of rostral Brodmann area 6 in mental-operation tasks: An integrative neuroimaging approach. *Cerebral Cortex* 12(11):1157–70. [TH]
- Hanakawa, T., Immisch, I., Toma, K., Dimyan, M. A., van Gelderen, P. & Hallett, M. (2003c) Functional properties of brain areas associated with motor execution and imagery. *Journal of Neurophysiology* 89(2):989–1002. [OD, TH]
- Haykin, S. (2001) *Kalman filtering and neural networks*. Wiley. [PG, aRG]
- Hein, A. & Held, R. (1961) A neural model for labile sensorimotor coordinations. *Biological Prototypes and Synthetic Systems* 1:71–74. [DMM]
- Heisenberg, M. & Wolf, R. (1988) Reafferent control of optomotor yaw torque in *Drosophila melanogaster*. *Journal of Comparative Physiology A* 163:373–88. [BW]
- Held, R. (1961) Exposure history as a factor in maintaining stability of perception and coordination. *Journal of Nervous and Mental Disease* 132:26–32. [DMM]
- Henriques, D. Y., Klier, E. M., Smith, M. A., Lowy, D. & Crawford, J. D. (1998) Gaze-centered remapping of remembered visual space in an open-loop pointing task. *Journal of Neuroscience* 18:1583–94. [OD]
- Hershberger, W. (1976) Efference copy, the closed-loop analogue of von Holst's efference copy. *Cybernetics Forum* 8:97–102. [JSJ]
- Holmes, G. (1917) The symptoms of acute cerebellar injuries due to gunshot injuries. *Brain* 40:461–535. [EMH]
- (1922) The Croonian lectures on the clinical symptoms of cerebellar disease and their interpretation. *Lancet* 2:111–15. [EMH]
- Hommel, B., Muesseler, J., Aschersleben, G. & Prinz, W. (2001) The theory of event coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences* 24:849–78. [RIS]
- Honda, M., Wise, S. P., Weeks, R. A., Deiber, M. P. & Hallett, M. (1998) Cortical areas with enhanced activation during object centred spatial information processing: A PET study. *Brain* 121:2145–58. [BT]
- Houk, J. C., Singh, S. P., Fischer, C. & Barto, A. (1990) An adaptive sensorimotor network inspired by the anatomy and physiology of the cerebellum. In: *Neural networks for control*, ed. W. T. Miller, R. S. Sutton & P. J. Werbos. MIT Press. [aRG]
- Hubbard, E. M., Altschuler, E. L., Gregory, R. L., Whip, E., Heard, P. & Ramachandran, V. S. (2000) Psychophysics and neuropsychology of the size-weight illusion. *Society for Neuroscience Abstracts* 26 (No.167): 4. [EMH]
- Hubbard, E. M., Altschuler, E. L. & Ramachandran, V. S. (in preparation) Size matters: Relative contribution of size vs. shape to the size-weight illusion. [EMH]
- Hubbard, T. L. (1996) Representational momentum, centripetal force, and curvilinear impetus. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22:1049–60. [MWi]
- Hubbard, T. L. & Bharucha, J. J. (1988) Judged displacement in apparent vertical and horizontal motion. *Perception and Psychophysics* 44:211–21. [MWi]
- Humphrey, N. K. (1993) *A history of the mind*. Vintage Books. [PG]
- Hurley, S. (forthcoming a) Active perception and perceiving action: The shared circuits hypothesis. In: *Perceptual experience*, ed. T. Gendler & J. Hawthorne. Oxford University Press. [rRG]
- Hurley, S. (forthcoming b) The shared circuits hypothesis: A unified functional architecture for control, imitation, and simulation. In: *Perspectives on Imitation: From mirror neurons to memes*, ed. S. Hurley & N. Chater. MIT Press. [rRG]
- Imamizu, H., Miyauchi, S., Tamada, T., Sasaki, Y., Takino, R., Putz, B., Yoshioke, T. & Kawato, M. (2000) Human cerebellar activity reflecting an acquired internal model of a new tool. *Nature* 403:192–95. [aRG]
- Ito, M. (1970) Neurophysiological aspects of the cerebellar motor control system. *International Journal of Neurology* 7:162–76. [arRG]
- (1984) *The cerebellum and neural control*. Raven Press. [arRG]
- (1993) Movement and thought: Identical control mechanisms by the cerebellum. *Trends in Neural Science* 16:448–50. [OD]
- Jeannerod, M. (1994) The representing brain: Neural correlates of motor intention and imagery. *Behavioral and Brain Sciences* 17(2):187–245. [aRG, NS]
- (1995) Mental imagery in the motor context. *Neuropsychologia* 33:1419–32. [aRG]
- (2001) Neural simulation of action: A unifying mechanism for motor cognition. *Neuroimage* 14:103–109. [aRG]
- Jeannerod, M. & Frak, V. (1999) Mental imaging of motor activity in humans. *Current Opinion in Neurobiology* 9:735–39. [aRG]
- Johnson, M. (1987) *The body in the mind*. University of Chicago Press. [aRG]
- Johnson, S. H. (2000a) Imagining the impossible: Intact motor representations in hemiplegics. *Neuroreport* 11:729–32. [aRG]
- (2000b) Thinking ahead: The case for motor imagery in prospective judgements of prehension. *Cognition* 74(2000):33–70. [aRG]
- Johnson, S. H., Rotte, M., Grafton, S. T., Hinrichs, H., Gazzaniga, M. S. & Heinze, H. J. (2002) Selective activation of a parietofrontal circuit during implicitly imagined prehension. *Neuroimage* 17:1693–704. [OD]
- Johnson-Laird, P. N. (1983) *Mental models*. Harvard University Press/Cambridge University Press. [aRG, NN]
- (2001) Mental models and deduction. *Trends in Cognitive Sciences* 5(10):434–42. [aRG, HW]
- Jordan, J. S. (1998) Recasting Dewey's critique of the reflex-arc concept via a theory of anticipatory consciousness: Implications for theories of perception. *New Ideas in Psychology* 16(3):165–87. [JSJ]
- (2000) The role of “control” in an embodied cognition. *Philosophical Psychology* 13:233–37. [JSJ]
- (2003) The embodiment of intentionality In: *Dynamical systems approaches to embodied cognition*, ed. W. Tschacher, pp. 201–27. Springer Verlag. [JSJ]
- Jordan, M. I., Rumelhart, D. E. (1992) Forward models: Supervised learning with a distal teacher. *Cognitive Science* 16:307–54. [EMH]
- Kalman, R. E. (1960) A new approach to linear filtering and prediction problems. *Journal of Basic Engineering* 82(D):35–45. [aRG]
- Kalman, R. & Bucy, R. S. (1961) New results in linear filtering and prediction theory. *Journal of Basic Engineering* 83(D):95–108. [aRG, VGo]
- Karni, A., Meyer, G., Jezzard, P., Adams, M. M., Turner, R. & Ungerleider, L. G. (1995) Functional MRI evidence for adult motor cortex plasticity during motor skill learning. *Nature* 14:155–58. [BT]
- Karni, A., Meyer, G., Rey-Hipolito, C., Jezzard, P., Adams, M. M., Turner, R. & Ungerleider, L. G. (1998) The acquisition of skilled motor performance: Fast

- and slow experience-driven changes in primary motor cortex. *Proceedings of the National Academy of Sciences USA* 95:861–68. [BT]
- Karniel, A. (2002) Three creatures named “forward model.” *Neural Networks* 15:305–307. [BW]
- Kawato, M. (1989) Adaptation and learning in control of voluntary movement by the central nervous system. *Advanced Robotics* 3(3):229–49. [rRG]
- (1990) Computational schemes and neural network models for formation and control of multijoint arm trajectories. In: *Neural networks for control*, ed. W. T. Miller, R. S. Sutton & P. J. Werbos. MIT Press. [rRG, EMH]
- (1997) Bidirectional theory approach to consciousness. In: *Cognition, computation and consciousness*, ed. M. Ito, Y. Miyashita & E. T. Rolls. Oxford University Press. [OD]
- (1999) Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology* 9:718–27. [aRG]
- Kawato, M., Furukawa, K. & Suzuki, R. (1987) A hierarchical neural network model for control and learning of voluntary movement. *Biological Cybernetics* 57:169–85. [aRG, JSJ]
- Kosslyn, S. M. (1994) *Image and brain*. MIT Press. [aRG]
- Kosslyn, S. M., Alpert, N. M., Thompson, W. L., Maljkovic, V., Weise, S. B., Chabris, C. F., Hamilton, S. E., Rauch, S. L. & Buonomano, F. S. (1993) Visual-mental imagery activates topographically-organized visual cortex: PET investigations. *Journal of Cognitive Neuroscience* 5:263–87. [aRG, BT]
- Kosslyn, S. M., Ball, T. M. & Reiser, B. J. (1978) Visual imagers preserve metric spatial transformation: Evidence from studies of images scanning. *Journal of Experimental Psychology: Human Perception and Performance* 4:47–60. [rRG, BT]
- Kosslyn, S. M. & Sussman, A. L. (1995) Roles of imagery in perception: Or, there is no such thing as immaculate perception. In: *The cognitive neurosciences*, ed. M. S. Gazzaniga, pp. 1035–42. MIT Press. [aRG, OW]
- Kosslyn, S. M., Thompson, W. L., Kim, I. J. & Alpert, N. M. (1995) Topographical representations of mental images in primary visual cortex. *Nature* 378:496–98. [aRG]
- Kosslyn, S. M., Thompson, W. L., Wraga, M. J. & Alpert, N. M. (2001) Imagining rotation by endogenous versus exogenous forces: Distinct neural mechanisms. *NeuroReport* 12:2519–25. [ADS, BT]
- Krakauer, J. W., Ghilardi, M.-F. & Ghez, C. (1999) Independent learning of internal models for kinematic and dynamic control of reaching. *Nature Neuroscience* 2(11):1026–31. [aRG]
- Kuo, A. (1995) An optimal control model for analyzing human postural balance. *IEEE Transactions on Biomedical Engineering* 42(1):87–101. [DMM]
- Lakatos, I. (1970) Falsification and the methodology of scientific research programmes. In: *Criticism and the growth of knowledge*, ed. I. Lakatos & A. Musgrave, pp. 91–195. Cambridge University Press. [EC]
- Lakoff, G. (1987) *Women, fire and dangerous things: What categories reveal about the mind*. The University of Chicago Press. [aRG, NN]
- Lakoff, G. & Johnson, M. (1999) *Philosophy in the flesh*. Basic Books. [aRG]
- Lamm, C., Windischberger, C., Leodolter, U., Moser, E. & Bauer, H. (2001) Evidence for premotor cortex activity during dynamic visuospatial imagery from single-trial functional magnetic resonance imaging and event-related slow cortical potentials. *NeuroImage* 14:268–83. [aRG]
- Lang, W., Cheyne, D., Hollinger, P., Gerschke, W. & Lindinger, G. (1996) Electric and magnetic fields of the brain accompanying internal simulation of movement. *Brain Research* 3:125–29. [BT]
- Langacker, R. W. (1987) *Foundations of cognitive grammar, vol. I*. Stanford University Press. [aRG]
- (1990) *Concept, image and symbol: The cognitive basis of grammar*. Mouton de Gruyter. [aRG]
- (1991) *Foundations of cognitive grammar, vol. II*. Stanford University Press. [aRG]
- (1999a) *Grammar and conceptualization*. Mouton de Gruyter. [aRG]
- (1999b) Viewing in cognition and Grammar. In: *Grammar and conceptualization*. (Cognitive Linguistics Research 14.) Mouton de Gruyter. [rRG]
- Lashley, K. S. (1951) The problem of serial order in behavior. In: *Cerebral mechanisms in behavior*, ed. L. A. Jeffress. Wiley. [MLL]
- Latash, M. L. (1993) *Control of human movement*. Human Kinetics. [MLL]
- Lewontin, R. C. (2001) *The triple helix: Gene, organism, and environment*. Harvard University Press. [HW]
- Liepmann, H. (1905) Die Linke Hemisphaere und das Handeln. *Muenchener Medizinische Wochenschrift* 48:2322–26, 49:2375–78. [BT]
- Llinas, R. & Pare, D. (1991) On dreaming and wakefulness. *Neuroscience* 44(3):521–35. [aRG]
- Lotze, M., Montoya, P., Erb, M., Hulsmann, E., Flor, H., Klose, U., Birbaumer, N. & Grodd, W. (1999) Activation of cortical and cerebellar motor areas during executed and imagined hand movements: An fMRI study. *Journal of Cognitive Neuroscience* 11:491–501. [BT]
- Luenberger, D. (1971) An introduction to observers. *IEEE Transactions on Automatic Control* 36(5):456–60. [DMM]
- Luria, A. R. (1973) *The working brain: An introduction to neuropsychology*. Penguin. [OW]
- Mach, E. (1896) *Contributions to the analysis of sensations*. Open Court. [aRG]
- Mahoney, M. J. & Avenier, M. (1987) Psychology of the elite athlete. An explorative study. *Cognitive Therapy and Research* 1:135–41. [NS]
- Maravita, A., Spence, C. & Driver, J. (2003) Multisensory integration and the body schema: Close to hand and within reach. *Current Biology* 13:(R)531–39. [BT]
- Marr, D. (1982) *Vision*. Freeman. [PG]
- Mataric, M. (1992) Integration of representation into goal-driven behavior-based robots. *IEEE Transactions on Robotics and Automation* 8(3):304–12. [LAS, GS]
- Matthews, P. B. C. (1959) The dependence of tension upon extension in the stretch reflex of the soleus of the decerebrate cat. *Journal of Physiology* 47:521–46. [MLL]
- Mehta, B. & Schaal, S. (2002) Forward models in visuomotor control. *Journal of Neurophysiology* 88(2):942–53. [aRG, MLL]
- Mel, B. W. (1986) A connectionist learning model for 3-d mental rotation, zoom, and pan. In: *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*, pp. 562–71. Erlbaum. [aRG]
- (1988) MURPHY: A robot that learns by doing. In: *Neural information processing systems*, ed. D. Z. Anderson. American Institute of Physics. [aRG]
- (1991) A connectionist model may shed light on neural mechanisms for visually guided reaching. *Journal of Cognitive Neuroscience* 3(3):273–92. [OW]
- Meltzoff, A. N. & Moore, K. M. (1977) Imitation of facial and manual gestures by human neonates. *Science* 198:75–78. [VS]
- Merfeld, D. M. (1995a) Modeling human vestibular responses during eccentric rotation and off vertical axis rotation. *Acta Oto-Laryngologica* (Supplement) 520:354–59. [DMM]
- (1995b) Modeling the vestibulo-ocular reflex of the squirrel monkey during eccentric rotation and roll tilt. *Experimental Brain Research* 106:123–34. [DMM]
- Merfeld, D. M., Young, L., Oman, C. & Shelhamer, M. (1993) A multi-dimensional model of the effect of gravity on the spatial orientation of the monkey. *Journal of Vestibular Research* 3:141–61. [DMM]
- Merfeld, D. M. & Zupan, L. H. (2002) Neural processing of gravito-inertial cues in humans. III. Modeling tilt and translation responses. *Journal of Neurophysiology* 87(2):819–33. [DMM]
- Merfeld, D. M., Zupan, L. H. & Peterka, R. (1999) Humans use internal models to estimate gravity and linear acceleration. *Nature* 398:615–18. [DMM]
- Miall, R. C. (1998) The cerebellum, predictive control and motor coordination. *Sensory Guidance of Movement, Novartis Foundation Symposium* 218:272–90. [EMH]
- Miall, R. C. & Wolpert, D. M. (1996) Forward models for physiological motor control. *Neural Networks* 9(8):1265–79. [TGC]
- Miles, C. F. & Rogers, D. (1993) A biologically motivated associative memory architecture. *International Journal of Neural Systems* 4(2):109–27. [aRG]
- Millikan, R. G. (1984) *Language, thought, and other biological categories*. MIT Press. [GS]
- (1993) *White Queen psychology and other essays for Alice*. MIT Press. [GS]
- Milner, A. D. & Goodale, M. A. (1996) *Visual brain in action*. Oxford University Press. [VG]
- Mohl, B. (1989) Short-term learning during flight control in *Locusta migratoria*. *Journal of Comparative Physiology* A163:803–12. [BW]
- Nair, D. G., Purcott, K. L., Fuchs, A., Steinberg, F. & Kelso, J. A. (2003) Cortical and cerebellar activity of the human brain during imagined and executed unimanual and bimanual action sequences: A functional MRI study. *Brain Research. Cognitive Brain Research* 15:250–60. [OD]
- Naito, E., Ehrsson, H. H., Geyer, S., Zilles, K. & Roland, P. E. (1999) Illusory arm movements activate cortical motor areas: A PET study. *Journal of Neuroscience* 19:6134–44. [NS]
- Naito, E., Kochiyama, T., Kitada, R., Nakamura, S., Matsumura, M., Yonekura, Y. & Sadato, N. (2002) Internally simulated movement sensations during motor imagery activate cortical motor areas and the cerebellum. *Journal of Neuroscience* 22:3683–91. [aRG, NS]
- Newton, N. (1996) *Foundations of understanding*. John Benjamins. [NN]
- Niedenthal, P. M., Barsalou, L. W., Winkielman, P., Krauth-Gruber, S. & Ric, F. (in press) Embodiment in attitudes, social perception, and emotion. *Personality and Social Psychology Review*. [CLR]
- Niedenthal, P. M., Brauer, M., Halberstadt, J. B. & Innes-Ker, A. H. (2001) When did her smile drop? Facial mimicry and the influences of emotional state on the detection of change in emotional expression. *Cognition and Emotion* 15:853–64. [CLR]
- Nolfi, S. & Tani, J. (1999) Extracting regularities in space and time through a cascade of prediction networks: The case of a mobile robot navigating in a structured environment. *Connection Science* 11(2):129–52. [aRG]
- Oman, C. (1982) A heuristic mathematical model for the dynamics of sensory

- conflict and motion sickness. *Acta Oto-Laryngologica* (Suppl.) 392:1–44. [DMM]
- (1990) Motion sickness: A synthesis and evaluation of the sensory conflict theory. *Canadian Journal of Physiology and Pharmacology* 68(2):294–303. [DMM]
- (1991) Sensory conflict in motion sickness: An observer theory approach. In: *Pictorial communication in virtual and real environments*, ed. S. Ellis, pp. 362–76. Taylor & Francis. [DMM]
- O'Regan, K. & Noë, A. (2001) A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences* 24(5):939–73. [GS]
- O'Reilly, R. & Munakata, Y. (2000) *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. MIT Press. [FCG]
- Ostry, D. J. & Feldman, A. G. (2003) A critical evaluation of the force control hypothesis in motor control. *Experimental Brain Research* 153:275–88. [RB, MLL]
- Palmer, S. (1978) Fundamental aspects of cognitive representation. In: *Cognition and categorization*, ed. E. Rosch & B. B. Lloyd. Erlbaum. [HW]
- Parsons, L. M. & Fox, P. T. (1998) The neural basis of implicit movements used in recognizing hand shape. *Cognitive Neuropsychology* 15:583–615. [BT]
- Pavani, F., Spence, C. & Driver, J. (2000) Visual capture of touch: Out-of-the-body experiences with rubber gloves. *Psychological Science* 11:353–59. [BT]
- Pellizzer, G., Sargent, P. & Georgopoulos, A. P. (1995) Motor cortical activity in a context-recall task. *Science* 269:702–705. [BT]
- Piaget J. (1947) *La psychologie de l'intelligence*. Armand Colin. [OW]
- Pisella, L., Grea, H., Tilikete, C., Vighetto, A., Desmurget, M., Rode, G., Boisson, D. & Rossetti, Y. (2000) An "automatic pilot" for the hand in human posterior parietal cortex: Toward reinterpreting optic ataxia. *Nature Neuroscience* 3(7):729–36. [VG]
- Poock, K. (1964) Phantoms following amputation in early childhood and in congenital absence of limbs. *Cortex* 1:269–75. [VS]
- Poock, K. & Orgass, B. (1971) The concept of the body schema: A critical review and some experimental results. *Cortex* 7(3):254–77. [VS]
- Porro, C. A., Francescato, M. P., Cettolo, V., Diamond, M. E., Baraldi, P., Zuiani, C., Bazzocchi, M. & di Prampero, P. E. (1996) Primary motor and sensory cortex activation during motor performance and motor imagery: A functional magnetic resonance imaging study. *Journal of Neuroscience* 16:7658–98. [BT]
- Poulet, J. F. A. & Hedwig, B. (2002) A corollary discharge maintains auditory sensitivity during sound production. *Nature* 418:872–76. [BW]
- Povinelli, D. J. (2000) *Folk physics for apes*. Oxford University Press. [PG]
- Powers, W. T. (1973) *Behavior: The control of perception*. Aldine. [JSJ]
- Pylyshyn, Z. W. (2001) Visual indexes, preconceptual objects, and situated vision. *Cognition* 80(1–2):127–58. [aRG]
- Quaia, C., Lefevre, P. & Optican, L. M. (1999) Model of the control of saccades by superior colliculus and cerebellum. *Journal of Neurophysiology* 82(2):999–1018. [TGC]
- Rao, R. P. N. & Ballard, D. H. (1999) Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience* 2(1):79–87. [aRG, VG]
- Reason, J. (1977) Learning to cope with atypical force environments. In: *Adult learning*, ed. M. Howe, pp. 203–22. Wiley. [DMM]
- (1978) Motion sickness adaptation: A neural mismatch model. *Journal of the Royal Society of Medicine* 71:819–29. [DMM]
- Reed, C. (2002a) Chronometric comparisons of imagery to action: Visualizing versus physically performing springboard dives. *Memory and Cognition* 30(8):1169–78. [OD, CLR]
- (2002b) What is the body schema? In: *The imitative mind*, ed. A. Meltzoff & W. Prinz, pp. 233–46. Cambridge University Press. [VS]
- Reed, C. L. & O'Brien, C. F. (1996) Motor imagery deficit in patients with Parkinson's Disease. Paper presented at the 3rd meeting of the Cognitive Neuroscience Society, San Francisco, 1996. [CLR]
- Reisberg, D. & Chambers, D. (1991) Neither pictures nor propositions: What can we learn from a mental image? *Canadian Journal of Psychology* 45:366–52. [MWI]
- Reisberg, D., Smith, J. D., Baxter, D. A. & Sonenshine, M. (1989) "Enacted" auditory images are ambiguous: "Pure" auditory images are not. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology* 41A:619–41. [MWI]
- Richter, W., Somorjai, R., Summers, R., Jarnasz, M., Menon, R. S., Gati, J. S., Georgopoulos, A. P., Tegeler, C., Ugurbil, K. & Kim, S. G. (2000) Motor area activity during mental rotation studied by time-resolved single-trial fMRI. *Journal of Cognitive Neuroscience* 12(2):310–20. [aRG, BT]
- Rizzolatti, G., Fadiga, L., Fogassi, L. & Gallese, V. (1999) Resonance behaviors and mirror neurons. *Archives Italiennes de Biologie* 137(2–3):85–100. [TGC]
- Roland, P. E., Skinhoj, E., Lassen, N. A. & Larsen, B. (1980) Different cortical areas in man in organization of voluntary movements in extrapersonal space. *Journal of Neurophysiology* 43:137–50. [BT]
- Ross, H. E. (1966) Sensory information necessary for the size-weight illusion. *Nature* 212:650. [EMH]
- Ross, H. E. & Gregory, R. L. (1970) Weight illusions and weight discrimination: A revised hypothesis. *Quarterly Journal of Experimental Psychology* 22:318–28. [EMH]
- Rozin, P. (1976) The evolution of intelligence and access to the cognitive unconscious. In: *Progress in psychobiology and physiological psychology, vol. 6*, ed. J. M. Sprague & A. N. Epstein, pp. 245–80. Academic Press. [OW]
- Rumelhart, D. E. & Norman, D. A. (1988) Representation in memory. In: *Stevens' handbook of experimental psychology*, ed. R. C. Atkinson, R. J. Herrnstein, G. Lindzey & R. D. Luce. Wiley. [HW]
- Rumiati, R. I., Tomasino, B., Vorano, L., Umiltà, C. & de Luca, G. (2001) Selective deficit of imagining finger configurations. *Cortex* 37:730–33. [BT]
- Ryle, G. (1949) *The concept of mind*. Barnes and Noble. [NN]
- Sadato, N., Campbell, G., Ibanez, V., Deiber, M. & Hallett, M. (1996) Complexity affects regional cerebral blood flow change during sequential finger movements. *Journal of Neuroscience* 16(8):2691–700. [TH]
- Sathian, K., Prather, S. C. & Zhang, M. (2004) Visual cortical involvement in normal tactile perception. In: *The handbook of multisensory processes*, ed. G. Calvert, C. Spence & B. Stein, pp. 703–709. MIT Press. [KS]
- Scholz, J. P., Schöner, G. & Latash, M. L. (2000) Identifying the control structure of multi-joint coordination during pistol shooting. *Experimental Brain Research* 135:382–404. [RB]
- Schroeder, C. E. & Foxe, J. J. (2002) The timing and laminar profile of converging inputs to multisensory areas of the macaque neocortex. *Cognitive Brain Research* 14:187–98. [KS]
- Schubotz, R. I. & von Cramon, D. Y. (2001) Functional organization of the lateral premotor cortex: fMRI reveals different regions activated by anticipation of object properties, location and speed. *Cognitive Brain Research* 11:97–112. [RIS]
- (2002) Predicting perceptual events activates corresponding motor schemes in lateral premotor cortex: An fMRI study. *NeuroImage* 15:787–96. [RIS]
- (2003) Functional-anatomical concepts of human premotor cortex: Evidence from fMRI and PET studies. *NeuroImage* 20:S120–S131. (Special Issue). [RIS]
- Schubotz, R. I., von Cramon, D. Y. & Lohmann, G. (2003) Auditory what, where, and when: A sensory somatotopy in lateral premotor cortex. *NeuroImage* 20:173–85. [RIS]
- Schwartz, D. L. (1999) Physical imagery: Kinematic versus dynamic models. *Cognitive Psychology* 38:433–64. [aRG]
- Schwobed, J., Boronat, C. B. & Coslett, H. B. (2002) The man who executed "imagined" movements: Evidence for dissociable components of the body schema. *Brain and Cognition* 50:1–16. [VG]
- Servos, P., Matin, L. & Goodale, M. A. (1995) Dissociation between two modes of spatial processing by a visual form agnostic. *NeuroReport* 6:1893–96. [VG, rRG]
- Servos, P., Osu, R., Santi, A. & Kawato, M. (2002) The neural substrates of biological motion perception: An fMRI study. *Cerebral Cortex* 12:772–82. [OD]
- Shadmehr, R. & Wise, S. P. (2003) Motor learning and memory for reaching and pointing. In: *The new cognitive neurosciences, 3rd edition*, ed. M. S. Gazzaniga. MIT Press. [OD]
- Shepard, R. N. & Metzler, J. (1971) Mental rotation of three-dimensional objects. *Science* 171:701–703. [TD]
- Siegal, M. & Varley, R. (2002) Neural systems involved in "theory of mind." *Nature Reviews. Neuroscience* 3:463–71. [OD]
- Sirigu, A. & Duhamel, J. R. (2001) Motor and visual imagery as two complementary but neurally dissociable mental processes. *Journal of Cognitive Neuroscience* 13:910–19. [VG, BT]
- Sirigu, A., Duhamel, J. R., Cohen, L., Pillon, B., Dubois, B. & Agid, Y. (1996) The mental representation of hand movements after parietal cortex damage. *Science* 273(5281):1564–68. [VG, BT]
- Smyrnis, N., Taira, M., Ashe, J. & Georgopoulos, A. P. (1992) Motor cortical activity in a memorized delay task. *Experimental Brain Research* 92:139–51. [BT]
- Snyder, L. H., Grieve, K. L., Brotchie, P. & Andersen, R. A. (1998) Separate body- and world-referenced representations of visual space in parietal cortex. *Nature* 394:887–91. [KS]
- Spencer, J. P. & Schöner, G. (2003) Bridging the representational gap in the dynamic systems approach to development. *Developmental Science* 6(4):392–412. [FCG]
- Sperry, R. W. (1950) Neural basis of the spontaneous optokinetic response produced by vision inversion. *Journal of Comparative and Physiological Psychology* 43:482–89. [DMM, BW]
- Stein, B. E. & Meredith, M. A. (1993) *Merging of the senses*. MIT Press. [KS]
- Stein, L. A. (1994) Imagination and situated cognition. *Journal of Experimental and Theoretical Artificial Intelligence* 6:393–407. [aRG, LAS]

- Stephan, K. M., Fink, G. R., Passingham, R. E., Silbersweig, D., Ceballos-Baumann, A. O., Frith, C. D. & Frackowiack, R. S. J. (1995) Functional anatomy of the mental representation of upper extremity movements in healthy subjects. *Journal of Neurophysiology* 73:373–86. [BT]
- Sternad, D. (2002) Wachholder, K & Altenburger, H (1927) Foundational experiments for current hypotheses on equilibrium point control in voluntary movements. *Motor Control* 6:299–318. [Historical overview, English translation, and commentaries on Wachholder & Altenburger 1927 by D. Sternad.] [MLL]
- Stojanov, G. (1997) Expectancy theory and interpretation of electroencephalograms (EXG) curves in the context of biological and machine intelligence. Ph.D. Thesis, Electrical Engineering Faculty, Saints Cyril and Methodius University, Skopje, Macedonia. [GS]
- Stojanov, G., Bozinovski, S. & Bozinovska, L. (1996) AV control system which makes use of environment stabilizations. In: *SPIE Proceedings, vol. 2903: Mobile Robots XI and Automated Vehicle Control Systems*, ed. C. H. Kenyon & P. Kachroo, pp. 44–51. SPIE. [GS]
- Stojanov, G., Bozinovski, S. & Trajkovski, G. (1997a) Interactionist expectative view on agency and learning. *IMACS Journal of Mathematics and Computers in Simulation* 44:295–310. [GS]
- (1997b) The status of representation in behaviour based robotic systems: The problem and a solution. Paper presented at the IEEE Conference on Systems, Man, and Cybernetics, Orlando, FL, 1997. [GS]
- Stojanov, G., Stefanovski, S. & Bozinovski, S. (1995) Expectancy based emergent environment models for autonomous agents. *Proceedings of the 5th International Symposium on Automatic Control and Computer Science, Iasi, Romania* 1:217–21. [GS]
- Sutton, R. S. & Barto, A. G. (1998) *Reinforcement learning: An introduction*. MIT Press. [HW]
- Tagaris, G. A., Richter, W., Kim, S.-G., Pellizzer, G. & Anderson, P. (1998) Functional magnetic resonance imaging of mental rotation and memory scanning: A multidimensional scaling analysis of brain activation patterns. *Brain Research Reviews* 26:106–12. [BT]
- Talmy, L. (2000a) Fictive motion in language and ‘ception’. In: *Toward a cognitive semantics, vols. 1 & 2*, ed. L. Talmy. MIT Press. [rRG]
- (2000b) *Toward a cognitive semantics*. MIT Press. [aRG]
- Tarsitano, M. S. & Andrew, R. (1999) Scanning and route selection in the jumping spider *Portia labiata*. *Animal Behaviour* 280:255–65. [BW]
- Thelen, E., Schöner, G., Scheier, C. & Smith, L. B. (2001) The dynamics of embodiment: A dynamic field theory of infant perseverative reaching errors. *Behavioral and Brain Sciences* 24:1–86. [FCG]
- Thomassen, A. J. W. M. & Tibosch, H. J. C. M. (1991) A quantitative model of graphic production. In: *Tutorials in motor neuroscience*, ed. G. E. Stelmach & J. Requin. Kluwer. [ADS]
- Todorov, E. & Jordan, M. I. (2002) Optimal feedback control as a theory of motor coordination. *Nature Neuroscience* 5:1226–35. [VG]
- Tomasello, M. (1999) *The cultural origins of human cognition*. Harvard University Press. [PG]
- Tomasino, B., Borroni, P., Isaja, A., Baldiserra, F. & Rumiati, R. I. (in press) The primary motor cortex subserves not only movements but also their imagination. *Cognitive Neuropsychology*. [BT]
- Turvey, M. T. (1990) Coordination. *American Psychologist* 45:938–53. [RB]
- Tversky, B. (2000) Remembering spaces. In: *The Oxford handbook of memory*, ed. E. Tulving & F. I. M. Craik. Oxford University Press. [HW]
- Ungerleider, L. G. & Haxby, J. V. (1994) “What” and “where” in the human brain. *Current Opinion in Neurobiology* 4(2):157–65. [aRG]
- van Beers, R. J., Sittig, A. C. & Gon, J. J. (1999) Integration of proprioceptive and visual position-information: An experimentally supported model. *Journal of Neurophysiology* 81:1355–64. [VG]
- van Beers, R. J., Wolpert, D. M. & Haggard, P. (2002) When feeling is more important than seeing in sensorimotor adaptation. *Current Biology* 12:834–37. [VG]
- van der Meulen, J. H. P., Goossens, R. H. J. M., van der Gon, J. J. D., Gielen, C. C. A. M. & Wilhelm, K. (1990) Mechanisms underlying accuracy in fast goal-directed arm movements in man. *Journal of Motor Behavior* 22(1):67–84. [aRG]
- van Galen, G. P. (1980) Handwriting and drawing: A two stage model of complex motor behaviour. In: *Tutorials in motor behaviour*, ed. G. E. Stelmach & J. Requin. North-Holland. [ADS]
- van Hoek, K. (1995) Conceptual reference points: A cognitive grammar account of pronominal anaphora constraints. *Language* 71(2):310–40. [aRG]
- (1997) *Anaphora and conceptual structure*. University of Chicago Press. [aRG]
- van Pabst, J. V. L. & Krekel, P. F. C. (1993) Multi sensor data fusion of points, line segments and surface segments in 3D space. In: *7th International Conference on Image Analysis and Processing, Capotolo, Monopoli, Italy*, pp. 174–82. World Scientific. [aRG]
- van Sommers, P. (1984) *Drawing and cognition*. Cambridge University Press. [ADS]
- (1989) A system for drawing and drawing-related neuropsychology. *Cognitive Neuropsychology* 6:117–64. [ADS]
- Vandervert, L. (1995) Chaos theory and the evolution of consciousness and mind: A thermodynamic-holographic resolution to the mind-body problem. *New Ideas in Psychology* 13(2):107–27. [JSJ]
- Verfaillie, K. & Daems, A. (2002) Representing and anticipating human actions in vision. *Visual Cognition* 9:217–32. [MWi]
- Verfaillie, K., de Troy, A. & van Rensbergen, J. (1994) Transsaccadic integration of biological motion. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20:649–70. [MWi]
- Verfaillie, K. & d’Ydewalle, G. (1991) Representational momentum and event course anticipation in the perception of implied periodical motions. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 17:302–13. [MWi]
- Vinter, A. (1994) Hierarchy among graphic production rules: A developmental approach. In: *Advances in handwriting and drawing: A multidisciplinary approach*, ed. C. Faure, P. Keuss, G. Lorette & A. Vinter. Europa. [ADS]
- Vinter, A. & Perruchet, P. (1999) Isolating unconscious influences: The neutral parameter procedure. *Quarterly Journal of Experimental Psychology* 52A:857–75. [ADS]
- Viviani, P. & Stucchi, N. (1989) The effect of movement velocity on form perception: Geometric illusions in dynamic displays. *Perception and Psychophysics* 46(3):266–74. [VG]
- (1992) Biological movements look uniform: Evidence of motor-perceptual interactions. *Journal of Experimental Psychology: Human Perception and Performance* 18(3):603–23. [VG]
- von Helmholtz, H. (1910) *Handbuch der physiologischen optik, vol. 3, 3rd edition*, ed. A. Gullstrand, J. von Kries & W. Nagel. Voss. [aRG]
- von Holst, E. (1954) Relations between the central nervous system and the peripheral organs. *British Journal of Animal Behavior* 2:89–94. [DMM, OW]
- von Holst, E. & Mittelstädt, H. (1950/1973) Das Refferenzprinzip: Wechselwirkungen zwischen Zentralnerven-system und Peripherie. *Naturwissenschaften* 37:467–76. (Original German publication, 1950.) English translation, 1973: The refference principle. In: *The behavioral physiology of animals and man. The collected papers of Erich von Holst*, trans. R. Martin, pp. 139–73. University of Miami Press. [JSJ, MLL, DMM, BW]
- von Uexküll, J. (1926) *Theoretische Biologie*. Suhrkamp. [DMM]
- Wachholder, K. & Altenburger, H. (1927/2002) Do our limbs have only one rest length. Simultaneously a contribution to the measurement of elastic forces in active and passive movements. *Pflüger’s Archive für die gesamte Physiologie* 215:627–40. (English translation by D. Sternad, 2002.) [see trans. in Sternad 2002]. [MLL]
- (2002) Foundational experiments for current hypotheses on equilibrium point control in voluntary movements. *Motor Control* 6:299–318. (English translation by D. Sternad, 2002.) [MLL]
- Walter, C. B., Swinnen, S. P., Dounskaia, N. & van Langendonk, H. (2001) Systematic error in the organization of physical action. *Cognitive Science* 25:393–422. [CBW]
- Wang, H., Johnson, T. R. & Zhang, J. (2001) The mind’s views of space. In: *Proceedings of the Third International Conference of Cognitive Science*. Beijing. [HW]
- Warren, W. H. Jr., Kay, B. A., Zosh, W. D., Duchon, A. P. & Sahuc, S. (2001) Optic flow is used to control human walking. *Nature Neuroscience* 4:213–16. [MLL]
- Weinstein, S. & Sersen, E. (1961) Phantoms in cases of congenital absence of limbs. *Neurology* 11:905–11. [VS]
- Weiss, Y., Simoncelli, E. P. & Adelson, E. H. (2002) Motion illusions as optimal percepts. *Nature Neuroscience* 5:598–604. [VG]
- Wellman, H. M. (1990) *The child’s theory of mind*. MIT Press. [aRG]
- Wexler, M. & Klam, F. (2001) Movement prediction and movement production. *Journal of Experimental Psychology: Human Perception and Performance* 27:48–64. [MW]
- Wexler, M., Kosslyn, S. M. & Berthoz, A. (1998) Motor processes in mental rotation. *Cognition* 68:77–94. [aRG, VG, MW]
- Wexler, M., Panerai, F., Lamouret, I. & Droulez, J. (2001) Self-motion and the perception of stationary objects. *Nature* 409:85–88. [MW]
- Wickens, T. D. (1993) Analysis of contingency tables with between-subjects variability. *Psychological Bulletin* 113:191–204. [ADS]
- Wiener, N. (1950) *The human use of human beings: Cybernetics and society*. Houghton Mifflin. [VG]
- Wiener, O. (1988) Form and content in thinking Turing machines. In: *The universal Turing machine*, ed. R. Herken, pp. 631–57. Oxford University Press. [OW]
- (1996) *Schriften zur Erkenntnistheorie*. Springer. [OW]
- (1998) “Klischee” als Bedingung intellektueller und künstlerischer Kreativität. In: *Literarische Aufsätze*, pp. 113–38. Löcker. [OW]

- (2000) Materialien zu meinem Buch "Vorstellungen." *Ausschnitt 5*, ed. F. Lesak. Technische Universität Wien. [OW]
- (2002) Anekdoten zu "Struktur." *Ausschnitt 7*, ed. F. Lesak, pp. 30–45. Technische Universität Wien. [OW]
- (forthcoming) *Vorstellungen*. Springer. [OW]
- Wilson, M. (2001) Perceiving imitatable stimuli: Consequences of isomorphism between input and output. *Psychological Bulletin* 127:543–53. [MWi]
- Wise, S. P., Moody, S. L., Blomstrom, K. J. & Mitz, A. R. (1998) Changes in motor cortical activity during visuomotor adaptation. *Experimental Brain Research* 121:285–99. [BT]
- Wohlschläger, A. (1998) Mental and manual rotation. *Journal of Experimental Psychology: Human Perception and Performance* 24:397–412. [MW]
- (2001) Mental object rotation and the planning of hand movements. *Perception and Psychophysics* 63:709–18. [ADS]
- Wolpert, D. M., Ghahramani, Z. & Flanagan, J. R. (2001) Perspectives and problems in motor learning. *Trends in Cognitive Sciences* 5(11):487–94. [aRG]
- Wolpert, D. M., Ghahramani, Z. & Jordan, M. I. (1995) An internal model for sensorimotor integration. *Science* 269:1880–82. [aRG, VG, EMH, MLL, CBW]
- Wolpert, D. M. & Kawato, M. (1998) Multiple paired forward and inverse models for motor control. *Neural Networks* 11(7–8):1317–29. [rRG]
- Wraga, M., Church, J. & Badre, D. (2002) Event-related fMRI study of imaginal self and object rotations. *Journal of Cognitive Neuroscience* 14:144. [ADS]
- Xing, J. & Andersen, R. A. (2000) Models of posterior parietal cortex which perform multimodal integration and represent space in several coordinate frames. *Journal of Cognitive Neuroscience* 12:601–14. [rRG]
- Yang, Y. & Bringsjord, S. (2003) Mental metalogic and its empirical justifications: The case of reasoning with quantifiers and predicates. *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society*, ed. R. Alterman & D. Kirsch, pp. 1275–80. Lawrence Erlbaum Associates. [HW]
- Yantis, S. (1992) Multielement visual tracking: Attention and perceptual organization. *Cognitive Psychology* 24(3):295–340. [aRG]
- Zacharias, G. L. & Young, L. R. (1981) Influence of combined visual and vestibular cues on human perception and control of horizontal rotation. *Experimental Brain Research* 41:159–71. [DMM]
- Zajac, F. E., Neptune, R. R. & Kautz, S. A. (2002) Biomechanics and muscle coordination of human walking. Part I: Introduction to concepts, power transfer, dynamics and simulations. *Gait and Posture* 16:215–32. [RB]
- Zajonc, R. B. & Markus, H. (1984) Affect and cognition: The hard interface. In: *Emotions, cognition, and behavior*, ed. C. Izard, J. Kagan & R. B. Zajonc, pp. 73–102. Cambridge University Press. [CLR]
- Zaretsky, M. & Rowell, C. H. F. (1979) Saccadic suppression by corollary discharge in the locust. *Nature* 280:583–85. [BW]
- Zhang, J. (1997) The nature of external representations in problem solving. *Cognitive Science* 21(2):179–217. [HW]
- Zhang, J. & Norman, D. A. (1994) Representations in distributed cognitive tasks. *Cognitive Science* 18:87–122. [HW]
- Zupan, L., Droulez, J., Darlot, C., Denise, P. & Maruani, A. (1994) Modelization of vestibulo-ocular reflex (VOR) and motion sickness prediction. Paper presented at the International Congress on Application of Neural Networks, Sorrento, Italy, 1994. [DMM]
- Zupan, L. H. & Merfeld, D. M. (2003) Neural processing of gravito-inertial cues in humans, IV. Influence of visual rotational cues during roll optokinetic stimuli. *Journal of Neurophysiology* 89(1):390–400. [DMM]
- Zupan, L., Merfeld, D. M. & Darlot, C. (2002) Using sensory weighting to model the influence of canal, otolith and visual cues on spatial orientation and eye movements. *Biological Cybernetics* 86:209–30. [DMM]
- Zupan, L., Peterka, R. & Merfeld, D. (2000) Neural processing of gravito-inertial cues in humans: I. Influence of the semicircular canals following post-rotatory tilt. *Journal of Neurophysiology* 84:2001–15. [DMM]