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-*

https://doi.org/10.1017/S0140525X04260095 Published online by Cambridge University Press

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]

[a]*Department of Biomedical Engineering, Johns Hopkins University, Baltimore, MD 21205;* [b]*Department 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 functional 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; Servos 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