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 & Perruchet 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 suggest 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 – Corridorness, 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.

https://doi.org/10.1017/S0140525X0443009X Published online by Cambridge University Press