

Modular motor learning

Chris Miall

An interesting theory of sensorimotor control has been recently extended and simulated. The simulation can learn to control an arm in several mutually exclusive 'contexts', situations where the arm carries one of four objects with different mechanical properties. It provides a good theoretical framework for testing biological motor systems.

We live and move in a complex environment. Each morning, without apparent effort, we can put on a heavy coat, pick up a briefcase, drive or cycle to work, then navigate around a computer screen all day long using a mouse, and later work out on the squash courts. Even if you do only some of these, think for a moment of the challenges facing our sensorimotor system. Not only do we control our bodies but we do so despite constraining clothes and added weights on our limbs; we treat all manner of mechanical devices as if they were extensions of ourselves; and we cope with complex visuomotor transformations between eye and hand. How can we do so many different things so well?

One answer is that we have the ability to adapt to new situations, and so modify the appropriate neural circuits. But a more intriguing answer is that we appear to treat each situation as different, and develop and switch between control circuits appropriate for each sensorimotor 'context'. The way the brain learns control under each new context without forgetting the old ones is being explored by computational theory.

An extended sensorimotor model

Haruno, Wolpert and Kawato [1] have recently extended their earlier work [2] on 'MOSAIC' (modular selection and identification for control) to the point where they now show full simulations of a model that can learn and operate in multiple different sensorimotor contexts. The key feature of their system is that it uses modular pairs of forward and inverse models, and a 'responsibility signal' calculated for each module (see Box 1). In each module, the forward model generates a prediction of the outcome of the motor commands being issued. If the prediction is confirmed by reafferent feedback, then that

pair of forward and inverse models must have been right for the current situation, and they should assume high responsibility for that action. Other modules would not be have been so appropriate; their predictions would be inaccurate, and so their responsibility should be low.

In fact, the complete system that Haruno *et al.* report has several elements that tie together previous models. There is an underlying feedback control loop that drives the actions when the internal models cannot. The feedback-error signal [3] is then used to drive learning in the internal models, weighted by the responsibility signal for each, so that a forward/inverse pair that is close to a desired controller for the current context is modified, whereas modules that have low responsibilities are hardly affected. This allows separate modules to develop for each control context, without overwriting existing modules [2]. The advantage of using feedback-error learning is that the goal of the algorithm is clear cut: it aims to reduce the error in the output of the controller to zero, by a gradient-descent learning rule. Thus, when or if the feedback-error signal is reduced to zero, the system has achieved an ideal controller. It is probably a more realistic, and more reliable, scheme than other identification algorithms which aim to produce an ideal inverse of the system they control [4].

Predictions

MOSAIC also uses responsibility *predictors*, which use feedforward signals to estimate the context and can thus bias the responsibility signal even before any action is made. This uses a simple 25-pixel neural network 'retina' to recognize shapes and learn the relationship between each shape cue and the object. It then uses a feedback update rule, based on the hidden Markov model, to smooth the likely transitions between contexts. For example, we rarely encounter objects whose physical properties instantly change, so prior knowledge about the very last context can be used to calculate the probability of the current context [5], and the transition between contexts is probabilistic. In this way, the visual input selects the appropriate control module

based on prior knowledge, and the evolving reafferent signals during the action either confirm that the context is the one originally assumed or signal that a shift in context has occurred.

Fitting the model in the brain

Haruno *et al.*'s simulation has many elegant aspects, and the demonstration of a working model is of course an important step. It is worth considering how this computational system might be achieved in the brain, and how we might test the model's predictions. For example, the responsibility predictor must use external signals, from vision, to influence the choice of control modules. These visuomotor relationships are learnt and probably involve dorsal premotor, ventral prefrontal and basal ganglia circuits [6]. So one might test whether predictive control, or updates of context estimates, require these areas.

Next, the motor command output of the whole MOSAIC system is a weighted sum, combining the output of the inverse models within each module, independently weighted by their responsibility signals. This summed output must descend to the motorneurons, either directly or via the descending cortico-spinal pathways. The responsibility signal thus selects between MOSAIC modules after 'softmax' normalization, calculated from the combined responsibilities of all modules. In turn, the error in each forward-model prediction needs to be normalized to all other forward-model errors, and these then modulate the output of each inverse model according to the reciprocal of the errors; small prediction errors mean high responsibility and vice versa. The feedback-error signal reflects the output of the underlying feedback loop, and in the original proposal [7] was carried to the cerebellum by complex spike activity. Could these various architectural demands constrain the circuitry enough to localize it? My own instinct is to assume that the forward models are held in the cerebellum [8], and interact with inverse models perhaps in the motor cortex; the authors' instinct is to put both forward and inverse models in the cerebellum.

Box 1. Models

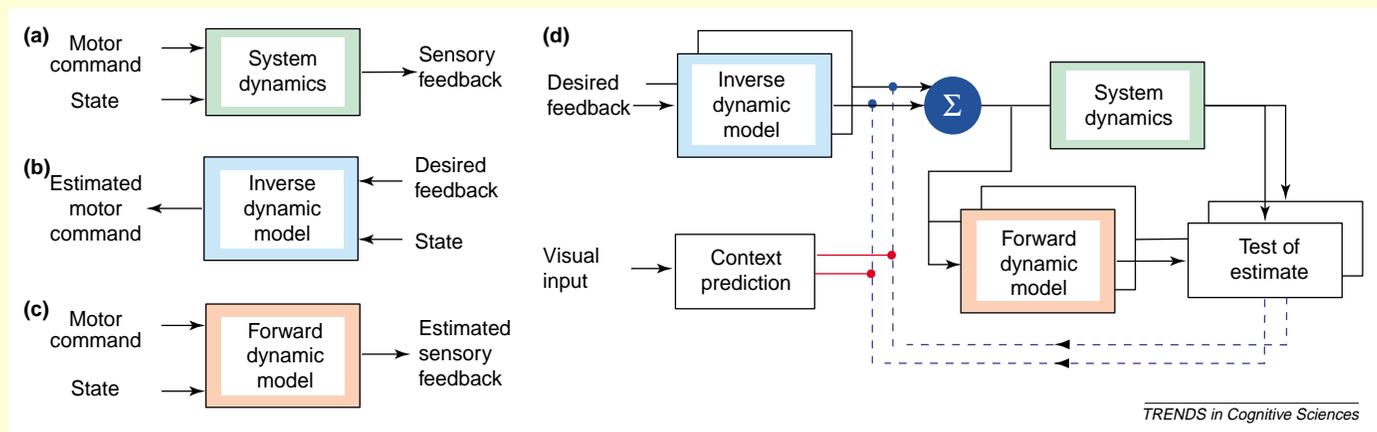
The motor system can be simplified into a 'black box' (Fig. 1a) that represents the dynamics of the spinal circuitry, skeletomuscular apparatus, and any additional objects that are being controlled, such as a squash racquet. Depending on the current state of this system, a motor command will cause it to move into a new state, detected by sensory channels (vision, proprioception), which might also contribute to the overall dynamics.

An 'inverse model' (Fig. 1b) describes the reverse of this control pathway: given the current state, it generates an estimate

of the motor command necessary to reach a desired sensory state. So the input is desired feedback or state, and the output is a motor command. By contrast, a 'forward model' (Fig. 1c) is an exact mimic of the motor system, and can generate an estimate of the sensory state that would be achieved if a motor command was followed.

The MOSAIC system uses multiple pairs of linked forward/inverse models (Fig. 1d; only 2 of each are shown here for simplicity). The inverse model outputs are weighted by current estimates of the

context (blue dashed lines) and summed together (blue circle) so that the most appropriate inverse model contributes most to the total motor command. An efferent copy of the combined motor command is then fed to each forward model that predicts the outcome. A reality check is performed, and the estimate of the context updated. Predictive visual cues can bias the context estimation before movement onset (red inputs). Other elements, such as an overall feedback loop and the training signals used to teach the models have not been shown.



TRENDS in Cognitive Sciences

Fig. 1. Models of the motor system. (See text for explanation.)

Questions to resolve

There are also important unanswered questions about the modular architecture that Haruno *et al.* propose. One is the extent of generalization they demonstrate. They have shown that if the system is trained to control four different physical objects (differing in mass, damping and spring constants), it can also control a novel one that is close to the centre of the 3-D space describing the original four. But if the new object is not within this space, the system cannot control it. So how fine-grained should we expect the representation of contexts to be? Do we need new modules for every different object (full coffee cup, half-empty coffee cup, etc.)? In fairness to the authors, we should not look to their work for the answer here, as it is likely that any biological system will tolerate some degree of inaccuracy in control, and compensate for the difference with corrective movements. Moreover, extended experience in a context will probably lead to a module with finer resolution. But it

does raise the question of when an existing module might be modified, rather than developing a new module.

Next, how does the nervous system assign resources to modules? We are born with a more or less blank slate, and learn new motor behaviours sequentially throughout life. So does the CNS have a rule to keep some modules uncommitted until they are eventually required, or does it reassign existing modules to new tasks? In either case, is there some limit to the number of separate modules it can hold?

This leads to the question of the extent to which MOSAIC modules can be combined. The most efficient outcome would be if the forward/inverse models could be linearly or non-linearly combined. For example, if a model of my arm could be combined with a module that coded in some way for another object, then could I combine the two models to predict how my arm behaves when I pick up the object? This is apparently not a trivial problem: the maths suggests that the forward models can be combined more easily than

the inverse models, but both would need to be combined for the scheme to work.

Finally, Haruno *et al.*'s simulations assume that there are negligible delays in the control pathways, but this is not true for biological systems because the sensory pathways, central neural computation, and efferent pathways all have significant delays. In MOSAIC, the internal predictions generated by the forward models are compared with the reafferent signal one computational time step ahead, but with physiologically delayed feedback this would introduce an error [8]. We can adapt to these errors – will MOSAIC?

References

- 1 Haruno, M., Wolpert, D.M. and Kawato, M. (2001) MOSAIC model for sensorimotor learning and control. *neural computation*. *Neural Comput.* 13, 2201–2220
- 2 Wolpert, D.M. and Kawato, M. (1998) Multiple paired forward and inverse models for motor control. *Neural Netw.* 11, 1317–1329
- 3 Kawato, M. (1990) Feedback-error-learning neural network for supervised motor learning. In: *Advanced Neural Computers* (Eckmiller, R., ed.), pp. 365–372, Elsevier

- 4 Karniel, A. *et al.* (2001) Best estimated inverse versus inverse of the best estimator. *Neural Netw.* 14, 1153–1159
- 5 Vetter, P. and Wolpert, D.M. (2000) Context estimation for sensorimotor control. *J Neurophysiol.* 84, 1026–1034
- 6 Jenmalm, P. and Johansson, R.S. (1997) Visual and somatosensory information about object

shape control manipulative fingertip forces.

J. Neurosci. 17, 4486–4499

- 7 Kawato, M. and Gomi, H. (1993) Feedback-error-learning model of cerebellar motor control.

In Role of the Cerebellum and Basal Ganglia in Voluntary Movement (Mano, N. *et al.*, eds), pp. 51–61, Elsevier

- 8 Miall, R.C. and Wolpert, D.M. (1996) Forward

models for physiological motor control. *Neural Netw.* 9, 1265–1279

Chris Miall

Dept of Physiology, University of Oxford, Parks Road, Oxford, UK OX1 3PT.

e-mail: chris.miall@physiol.ox.ac.uk

Are psychology's tribes ready to form a nation?

Daniel Gilbert

Collaboration between social psychologists and cognitive neuroscientists is giving rise to a new approach that its practitioners call 'social cognitive neuroscience'. Scientists from each discipline are using the theories and techniques of the other to generate new answers to fundamental questions about attitudes, beliefs, the self, moral judgment, and other issues. Is this interdisciplinary endeavor an exercise in wishful thinking and good intentions, or is it a preview of psychology's future?

Several decades ago, an eminent psychologist defined the field of psychology as 'a bunch of men standing on piles of their own crap, waving their hands and yelling "Look at me, look at me!"' Fortunately, things have changed quite a bit over the years, and the field is no longer composed entirely of men. The criticism is overstated, of course, but it does highlight one of psychology's most troubling shortcomings, namely, that psychologists often ignore work outside their own laboratories, usually ignore work outside their own sub-specialties, and almost always ignore work outside their own discipline. This parochialism is especially pronounced during transitional moments in the field's evolution, when the excitement generated by new ideas and new technologies seems to justify the sweeping away of history.

The emergence of cognitive neuroscience was one of the signal events in 20th century psychology, and psychologists have good reason to be optimistic about its future. How brains make minds is the critical missing piece in psychology's analysis of human behavior, and during the next few decades, cognitive neuroscience is sure to produce many useful insights and perhaps a few stunning ones. Alas, if this new enterprise is anything like its ancestors, its early impulse will be to invent itself by ignoring as much of the rest

of psychology as it can get away with, and there is already some evidence of this. Descartes made many errors, but failing to read his peers was not among them.

Given psychology's tendency to start each revolution from scratch, it is heartening to note that some researchers are making a concerted effort to ensure that cognitive neuroscience does not make the same mistake. In a recent article, Kevin Ochsner (a cognitive neuroscientist at Stanford University, CA, USA) and Matthew Lieberman (a social psychologist at UCLA, Los Angeles, CA, USA) have issued a clarion call for the integration of the neurological, cognitive and social levels of analysis [1]. Like most clarion calls, theirs is full of good intentions. Unlike most clarion calls, it is also full of good ideas about how to carry out that mission, and full of evidence that the integration is already underway. Ochsner and Lieberman review several problems on which cognitive neuroscientists and social psychologists are now working together, for example, the role of amygdala activation in stereotyping, hemispheric asymmetries and self-knowledge, amnesia and attitude change, and the role of the lateral fusiform gyrus in dispositional attribution. In each instance, Ochsner and Lieberman demonstrate how the two fields are collaborating, converging, and informing one another. The reason for this mutual attraction is obvious: cognitive neuroscience offers a new set of tools with which to examine enduring problems and holds out the prospect of grounding behavior in biology, whilst social psychology offers a treasure trove of theory and data about the kinds of problems our social brains were evolved to solve, and the kinds of solutions they have actually generated. The fruits of this social-cognitive neuroscience approach are already clear: Articles have appeared

in leading journals, conferences on social cognitive neuroscience (most notably those sponsored by Dartmouth University and UCLA) have attracted bright young people and well-established leaders from both disciplines, and federal granting agencies are paying the kind of attention that counts.

As with any marriage of true minds, this one admits of impediments, and these have mainly to do with the misgivings and misunderstandings that naturally arise whenever different tribes meet at the watering hole. In the privacy of their laboratories, social psychologists often marvel at the naïveté of neuroscientific research on 'social cognition', which all too often assumes that anything that hasn't been studied in a scanner hasn't been studied at all. Cognitive neuroscientists are similarly likely to roll their eyes at the naïveté of social psychologists, who happily (or haplessly) develop mentalistic theories without stopping to ask whether the 'machine' is actually capable of running the software. All of this may be true, but Ochsner and Lieberman have shown that some scientists have set aside their tribal prejudices long enough to recognize that although both disciplines can get along just fine without the other, both are enhanced when they do what they do best in each other's company. If Ochsner and Lieberman are right, psychologists might someday find themselves standing atop one giant heap, yelling 'Look at us! Look at us!'

Reference

- 1 Ochsner, K.N. and Lieberman, M.D. (2001) The emergence of social cognitive neuroscience. *Am. Psychol.* 56, 717–734

Daniel Gilbert

Dept of Psychology, William James Hall, Harvard University, Cambridge, MA 02138, USA.

e-mail: dtg@wjh.harvard.edu