

results from tests of this material in electrolyte cells have shown an acceptable lifetime, with a competitive energy per unit weight at reasonable rates of charge and discharge. Based on material costs, it is anticipated that the system might fall between aqueous batteries and lithium-ion batteries in terms of price.

As with any new technology, many tests must be done to fully characterize the system and to optimize the battery's behaviour, and these will no doubt be carried out by battery manufacturers with an interest in the system.

In the end, of course, it must eventually prove itself in the marketplace. But our need to find viable alternatives to hydrocarbon fuels means there is every incentive to make full and fair evaluations of technical advances such as this one. ■

George E. Blomgren is at Blomgren Consulting Services Ltd, 1554 Clarence Avenue, Lakewood, Ohio 44107, USA.

e-mail: geblomgren@prodigy.net

1. Aurbach, D. *et al.* *Nature* **407**, 724–727 (2000).
2. Connor, J., Reid, W. & Wood, G. J. *Electrochem. Soc.* **104**, 38–41 (1957).

Computational neuroscience

Building blocks of movement

Zoubin Ghahramani

When we walk, turn the pages of a book, or catch a ball, our brain is solving a 'control problem' — that of coordinating the activities of all the muscles required to achieve these voluntary movements. Roboticists have long appreciated the computational challenge of controlling such movements. For even a simplified description of the human arm, several pages of equations are required to express the forces that muscles need to exert to accelerate the arm in a given direction. And these equations would change when the arm picks up an object or moves in a different medium,

such as water, or as muscles tire. The challenge for neuroscientists has been to discover how these equations are represented in the brain, and how these representations adapt to suit new environments or tasks. On page 742 of this issue, Thoroughman and Shadmehr¹ provide some answers.

They do this by expanding on an idea that has captured the imagination of physiologists, modellers and psychologists alike. This idea is that the central nervous system creates seemingly complex behaviours by combining a relatively small number of simpler 'motor primitives'. Like a child putting

together simple building blocks to create elaborate structures, the brain — so the theory goes — puts together these motor primitives to create new behaviours.

What might constitute these 'building blocks' of movement? Many (not mutually exclusive) suggestions have been made, such as synergistic muscle contractions², elementary stable postures³ and simple movement strokes⁴. But the problem for experimentalists has been to establish the existence of these primitives, to obtain a 'picture' of what they look like, and to understand their role in everyday behaviours. Thoroughman and Shadmehr¹ now provide a mathematical description of the motor primitives that the central nervous system uses to learn to control the arm.

To do this, the authors make use of a well-studied⁵ task in which subjects make reaching movements while holding on to a handle attached to the end of a lightweight robotic arm (Fig. 1a). When the robotic arm is turned off, this is a very easy and natural task. When the robotic arm is turned on, however, it generates forces that simulate a very strange kind of environment, not previously experienced by the subjects. In this environment, a reaching movement in any direction by the subject causes forces perpendicular to the direction of movement and proportional to the velocity of movement, transforming previously straight movements into roughly spiral-shaped ones. It is well known that subjects eventually adapt to this type of force field, but Thoroughman and Shadmehr study this adaptation in detail. Their ingenious analysis reveals remarkable structure in the pattern of adaptation, giving us a glimpse of motor primitives in action.

The authors reason that if a reaching movement occurring within a force field causes the brain to learn, and if this learning is represented by changes to motor primitives, then changes will be seen in subsequent reaching movements — not only in the target direction but also in other directions. So, by working out how learning 'spills over', or generalizes, between movement directions, the mathematical shape of the motor primitives can be uncovered.

Indeed, Thoroughman and Shadmehr find that individual movements do have an effect on subsequent movements in other directions. This generalization decays with the angular distance between movement directions: learning how to move the arm in one direction results in partial learning of how to move the arm in nearby directions, but unlearning of movements in the opposite direction. The pattern of generalization suggests that the motor primitives map the desired velocity of the reaching hand to the force required to move the hand at this velocity. The pattern also suggests that mathematically the primitives have a gaussian shape (Fig. 1b).

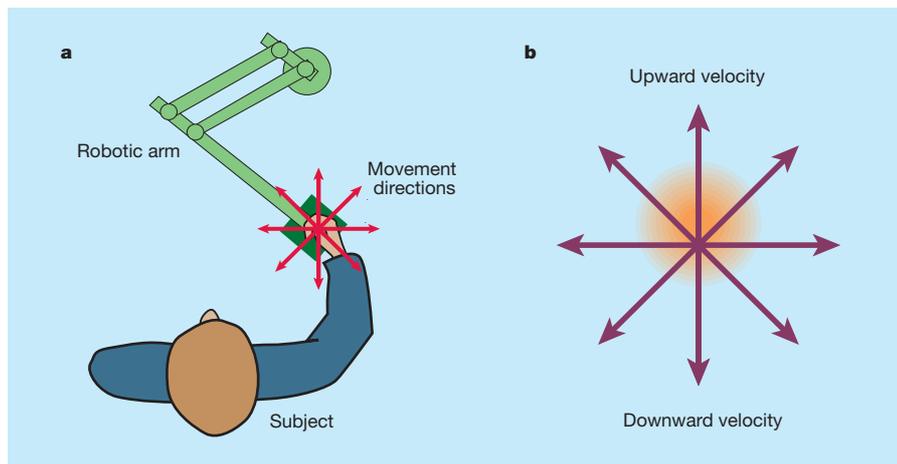


Figure 1 Producing complex movements from simple building blocks. Thoroughman and Shadmehr¹ investigated 'motor primitives', thought to represent the neural building blocks of complicated movements. a, The experiment. Human subjects held on to a robotic arm and moved their hand towards targets in eight different directions. The robotic arm exerted forces on the hand that depended on the velocity of movement of the hand. At first these forces affected the reaching movements by the subjects, but eventually the subjects adapted to the forces, allowing them to reach in the correct direction. b, The model. The pattern by which a reaching movement in the force field in one direction affected subsequent movements in other directions suggested a model in which the dynamics of movement are learned by changing and combining motor primitives that encode the force required to move the arm at a particular velocity. The primitives overlap (only one is shown, in orange), and have a broad gaussian tuning to desired hand velocity. The primitive shown is tuned to produce the force needed for movement in the upward direction, but spills over into other directions in a manner that matches the experimental results¹.

e-mail: zoubin@gatsby.ucl.ac.uk

1. Thoroughman, K. A. & Shadmehr, R. *Nature* **407**, 742–747 (2000).
2. Bernstein, N. *The Co-ordination and Regulation of Movements* (Pergamon, Oxford, 1967).
3. Bizzi, E., Mussa-Ivaldi, F. A. & Giszter, S. *Science* **253**, 287–291 (1991).
4. Morasso, P. & Mussa-Ivaldi, F. A. *Biol. Cybern.* **45**, 131–142 (1982).
5. Shadmehr, R. & Mussa-Ivaldi, F. A. *J. Neurosci.* **14**, 3208–3224 (1994).
6. Loeb, G. E. in *Progress in Brain Research* vol. 123 (ed. Binder, M. D.) 405–409 (Elsevier, Cambridge, 1999).
7. Todorov, E. *Nature Neurosci.* **3**, 391–398 (2000).
8. Goodbody, S. J. & Wolpert, D. M. *J. Neurophysiol.* **79**, 1825–1838 (1998).
9. Gandolfo, F., Li, C.-S. R., Benda, B. J., Padoa Schioppa, C. & Bizzi, E. *Proc. Natl Acad. Sci. USA* **97**, 2259–2263 (2000).

The authors then devised an adaptive, computational model composed of such primitives, and compared the behaviour of humans and the model in several new variations of the task. The model correctly predicted the S-shaped trajectories of human reaching movements; that these S-shaped trajectories would become straight if the proportion of ‘catch’ trials (when the robot was surreptitiously turned off) were decreased; and how subjects would adapt incorrectly to force fields that changed direction too rapidly.

These results are interesting for three reasons. First, it has been suggested⁶ that neither the inputs (‘proprioceptive’ neuronal signals) nor the outputs (muscle activations) of the human movement-control system show a simple relationship to external forces or the speed of movement of a reaching hand. So it is remarkable that Thoroughman and Shadmehr’s model — in which motor primitives map desired hand velocity to force — can account for the details of the time course and the end product of the human subjects’ adaptation. This type of mapping is exactly what makes control through motor primitives computationally attractive. The neural ‘controller’, presumably in the brain and spinal cord, need not concern itself with how hand velocities are computed or how forces are produced; its task is simply to learn how to map between the two.

All the same, we should not jump to conclusions from these appealing results. Motor primitives that appear to encode extrinsic variables, such as force and hand velocity, are somewhat divorced from the nitty-gritty of the muscle–skeletal system. Such primitives might turn out to be encoded instead in terms of more fundamental variables, such as muscle activations⁷.

Second, we already knew that the computations required to adapt to new, dynamic environments need to be implemented somewhere in the central nervous system. Thoroughman and Shadmehr’s results implicate one brain region — the cerebellum — as the area responsible. In particular, their finding that the motor primitives are velocity-tuned fits in with the encoding of velocity seen in Purkinje cells in the cerebellum. Other evidence also points to a key role for the cerebellum in motor learning. But of course things are not clear cut — previous studies using a different generalization paradigm⁸ and primate neurophysiology⁹ indicated that the brain region responsible for learning new movement dynamics could be the primary motor cortex. It is likely that changes in both cerebellum and primary motor cortex are involved in learning this particular task. Besides which, there is little reason to think that a computational element such as a motor primitive should map simply to brain structure.

Finally, these results are interesting from

a robotics perspective. Most artificial learning systems sluggishly accumulate small changes. But in humans, a single movement within the force field can have significant effects on subsequent movements — this biological movement-control system remains stable while adapting rapidly. By studying the building blocks of movement we not only learn more about the neurobiology of movement in humans. Some day, we may want to put those building blocks into robots, too. ■

Zoubin Ghahramani is in the Gatsby Computational Neuroscience Unit, University College London, London WC1N 3AR, UK.

Conservation biology

Seeds of doubt

Peter D. Moore

What is a weed to a farmer may well be a cherished wild flower to someone else, or an essential ecological component of a local flora. Hence the attempts, in response to diminishing botanical diversity, to sustain some wild plants by sowing their seeds in areas where they are under threat. But as Keller and colleagues show in *Journal of Applied Ecology*¹, this kind of conservation work can have adverse effects by diluting and modifying the local genetic resources of residual weed populations.

Agricultural weed control methods have generally met with considerable success, with the result that crops are more productive and harvests are less contaminated with unwanted seeds. But the side effects of weed decline include changed rural landscapes, falling numbers of invertebrates and even decreases in the populations of some birds. Weed seed mixtures are now commercially available (marketed, of course, as wild flowers), and an understandable response of environmentalists has been the deliberate dispersal of such seeds, especially along road margins in Europe and North America.

On the face of it, the re-establishment of more biodiverse agricultural landscapes is a creditable and straightforward aim. Simply allowing nature to take its course, by stopping herbicide treatment or simple disturbance of soil, or both, might lead to weed resurgence, but only of those that are well represented in the soil seedbank. Some weeds have seeds with a limited capacity to survive in soil, and these species in particular are candidates for supplementation with seeds from elsewhere. One disadvantage with this approach is that introduced plants, even those of the same species as local plants, might be poorly adapted to local conditions, including both the physical conditions of the habitat and the indigenous grazers, polli-

nators and pathogens. Can maladaptation of this sort be transferred to the residual local populations of plants?

To test this possibility, Keller *et al.*¹ studied three weed species, the common poppy (*Papaver rhoeas*), corncockle (*Agrostemma githago*; Fig. 1) and white campion (*Silene alba*). Their seeds are available commercially from several countries, including Switzerland, Germany, Britain, Hungary and the United States, and all are derived from local wild populations. The authors conducted a range of crossing experiments (always using Swiss material as mother plants) and measured such features as biomass yield, sur-

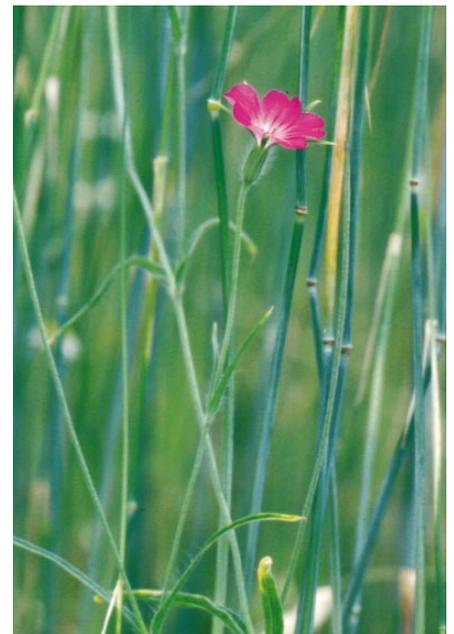


Figure 1 The corncockle. This weed (or flower, according to your point of view) is now rare in Britain but is still relatively common in southern and eastern Europe.

HOLT/ROB GIBBONS