

weather station at Grand Rapids, MI, located 35 km south of Grant. Soil measurements were made using Hobo data loggers (H08-031-08) fitted with probes (H08-031-08) buried 1.5 cm in the soil, the mean depth that *Rhagoletis* pupae overwinter²⁴.

Computer modelling of selection

Computer simulations were conducted using a discrete generation model mirroring the univoltine life-cycle of *Rhagoletis*. A random sample of 6% of the adult population enclosing under apple and haw trees was assumed to move to the alternate host each generation, matching gene flow estimates from mark-recapture studies⁹. Hard selection was invoked by weighting the proportions of immigrants and residents on a host by the mean fitness of the emigrant and resident race, respectively. After migration, adults randomly mated within host demes, with offspring experiencing viability selection based on absolute fitness values derived from the 26 °C (apple) or 22 °C (haw) prewinter, 30-week winter treatments. Given these values, the simulations converged on a stable equilibrium for each allozyme that maintained polymorphism and host-related differentiation regardless of initial allele frequencies.

Received 24 March; accepted 31 July 2000.

1. Darwin, C. *The Origin of Species by Means of Natural Selection* (Murray, London, 1859).
2. Dobzhansky, T. *Genetics and the Origins of Species* (Columbia Univ. Press, New York, 1937).
3. Mayr, E. *Systematics and the Origin of Species* (Columbia Univ. Press, New York, 1942).
4. Bush, G. L. in *Evolutionary Strategies of Parasitic Insects and Mites* (ed. Price, P. W.) 187–206 (Plenum, New York, 1975).
5. Bush, G. L. *The Taxonomy, Cytology and Evolution of the Genus Rhagoletis in North America* (Museum of Comparative Zoology, Cambridge, Massachusetts, 1966).
6. Prokopy, R. J., Diehl, S. R. & Cooley, S. S. Behavioral evidence for host races in *Rhagoletis pomonella* flies. *Oecologia* **76**, 138–147 (1988).
7. Feder, J. L., Chilcote, C. A. & Bush, G. L. Genetic differentiation between sympatric host races of *Rhagoletis pomonella*. *Nature* **336**, 61–64 (1988).
8. McPheron, B. A., Smith, D. C. & Berlocher, S. H. Genetic differences between *Rhagoletis pomonella* host races. *Nature* **336**, 64–66 (1988).
9. Feder, J. L. *et al.* Host fidelity is an effective pre-mating barrier between sympatric races of the apple maggot fly. *Proc. Natl Acad. Sci. USA* **91**, 7990–7994 (1994).
10. Feder, J. L. & Filchak, K. E. It's about time: The evidence for host plant-mediated selection in the apple maggot fly, *Rhagoletis pomonella*, and its implications for fitness trade-offs in phytophagous insects. *Ent. Exp. Appl.* **91**, 211–225 (1999).
11. Prokopy, R. J., Bennett, E. W. & Bush, G. L. Mating behavior in *Rhagoletis pomonella* (Diptera: Tephritidae). I. Site of assembly. *Canad. Ent.* **103**, 1405–1409 (1971).
12. Feder, J. L. & Bush, G. L. A field test of differential host plant usage between two sibling species *Rhagoletis pomonella* fruit flies (Diptera: Tephritidae) and its consequences for sympatric models of speciation. *Evolution* **43**, 1813–1819 (1989).
13. Reid, J. A. & Laing, J. E. Development threshold and degree-days to adult emergence for overwintering pupae of the apple maggot, *Rhagoletis pomonella* (Walsh) collected in Ontario. *Proc. Ent. Soc. Ontario* **197**, 19–22 (1976).
14. Via, S. in *Ecological Genetics* (ed. Real, L. A.) 58–85 (Princeton Univ. Press, Princeton, New Jersey, 1994).
15. Craig, T. P., Horner, J. D. & Itami, J. K. Hybridization studies on the host races of *Eurosta solidaginis*: Implications for sympatric speciation. *Evolution* **51**, 1552–1560 (1997).
16. Thompson, J. N. *The Coevolutionary Process* (Univ. Chicago Press, Chicago, 1994).
17. Abrahamson, W. G. *et al.* in *Gall-Forming Insects* (eds Price, P., Mattson, W. & Baranchilov, Y.) 208–222 (USDA Forest Service Tech. Report NC-174, St. Paul, Minnesota, 1994).
18. Wood, T. K. & Keese, M. C. Host-plant-induced assortative mating in *Echenopa* treehoppers. *Evolution* **44**, 619–628 (1990).
19. Tauber, C. A. & Tauber, M. J. in *Speciation and Its Consequences* (eds Otte, D. & Endler, J. A.) 307–344 (Sinauer, Sunderland, Massachusetts, 1989).
20. Fry, J. D. The evolution of host specialization: are trade-offs overrated? *Am. Nat.* **148**, S84–S107 (1996).
21. Feder, J. L. in *Endless Forms: Species and Speciation* (eds Howard, D. & Berlocher, S. H.) 130–144 (Oxford Univ. Press, New York, 1998).
22. Farrell, B. D. “Inordinate fondness” explained: Why are there so many beetles? *Science* **281**, 555–558 (1998).
23. Roethele, J. B. *et al.* Towards a molecular genetic linkage map for the apple maggot fly, *Rhagoletis pomonella* (Diptera:Tephritidae): a comparison of alternative strategies. *Ann. Entomol. Soc. Am.* **90**, 470–479 (1997).
24. Lathrop, F. H. & Nickels, C. B. *The biology and control of the blueberry maggot in Washington County Maine* (Tech. Bull. 275 US Dept of Agriculture, Washington DC, 1932).
25. Feder, J. L. The effects of parasitoids on sympatric host races of the apple Maggot fly, *Rhagoletis pomonella* (Diptera: Tephritidae). *Ecology* **76**, 801–813 (1995).
26. Filchak, K. E. *et al.* A field test for host-plant dependent selection on larvae of the apple maggot fly, *Rhagoletis pomonella*. *Evolution* **53**, 187–220 (1999).

Supplementary information is available on Nature’s World-Wide Web site (<http://www.nature.com>) or as paper copy from the London editorial office of Nature.

Acknowledgements

We thank S. Berlocher, J. Carbol, H. Dambroski, G. Dwyer, V. Gavrilovic, D. Gibbons, D. Lodge, B. Perry, D. Prokrym, U. Stolz, S. Via and the USDA lab at Niles, Michigan. This work was supported by grants from NSF (Young investigators, Integrative Research Challenges & GRT) and the USDA to J.L.F. and the University of Notre Dame.

Correspondence and requests for materials should be addressed to K.E.F. (e-mail: Filchak.1@nd.edu).

Learning of action through adaptive combination of motor primitives

Kurt A. Thoroughman*† & Reza Shadmehr*

* Department of Biomedical Engineering, Johns Hopkins University, 419 Traylor Building, 720 Rutland Avenue, Baltimore, Maryland 21205, USA

Understanding how the brain constructs movements remains a fundamental challenge in neuroscience. The brain may control complex movements through flexible combination of motor primitives¹, where each primitive is an element of computation in the sensorimotor map that transforms desired limb trajectories into motor commands. Theoretical studies have shown that a system’s ability to learn action depends on the shape of its primitives². Using a time-series analysis of error patterns, here we show that humans learn the dynamics of reaching movements through a flexible combination of primitives that have gaussian-like tuning functions encoding hand velocity. The wide tuning of the inferred primitives predicts limitations on the brain’s ability to represent viscous dynamics. We find close agreement between the predicted limitations and the subjects’ adaptation to new force fields. The mathematical properties of the derived primitives resemble the tuning curves of Purkinje cells in the cerebellum. The activity of these cells may encode primitives that underlie the learning of dynamics.

Studies of reaching movements have demonstrated that humans construct motor commands based on a prediction of forces that will be experienced in the upcoming movement³. When new forces are imposed on the arm, the prediction is in error and the arm does not follow the desired trajectory^{3,4}. With practice the motor commands are modified⁵ and the trajectory approximates the desired path. The learning of dynamics, however, affects movements outside the region of training^{3,6–8}, suggesting that the brain builds a state-dependent approximation of external forces⁹, called an internal model. Occasional movements with unexpectedly altered dynamics, termed ‘catch trials’, have been used to quantify how the internal model generalizes^{3,4}. Catch trials, however, not only test the internal model for a given movement but cause errors that in turn change the internal model and affect future movements. We demonstrate that the effect of errors experienced in a given movement on subsequent movements can reveal characteristics of primitives with which motor commands are generated.

We consider the internal model to be a sensorimotor map transforming desired arm trajectories into muscle forces^{10–12} through a flexible combination of a set of primitives:

$$\hat{\mathbf{f}} = W^T \mathbf{g}(\mathbf{x}^*, \dot{\mathbf{x}}^*, \ddot{\mathbf{x}}^*) \tag{1}$$

where T is the transpose operator, $\hat{\mathbf{f}}$ is a vector approximation of forces \mathbf{f} to be produced by muscles to compensate for task dynamics, and \mathbf{g} is a vector of scalar-valued primitives $[g_1, \dots, g_j]^T$. Although in general \mathbf{g} can depend on desired position, velocity and acceleration $(\mathbf{x}^*, \dot{\mathbf{x}}^*, \ddot{\mathbf{x}}^*)$, here we investigated learning of viscous forces and therefore considered a simpler subset of primitive functions that depended only on desired velocity. The internal model is learned through experience-dependent modification of the weight matrix W . Assuming a learning rule that minimizes $\tilde{\mathbf{f}}^2 \equiv \|\mathbf{f} - \hat{\mathbf{f}}\|^2$, W is adjusted after a movement (indexed 1) according to:

$$\Delta W_1 = -\eta \mathbf{g}(\dot{\mathbf{x}}_1^*) \tilde{\mathbf{f}}^T \tag{2}$$

where η is a constant learning step. This adaptation changes the

† Present address: Volen Center of Complex Systems and Department of Biology, Brandeis University, Waltham, Massachusetts 02454, USA.

internal model output in the subsequent movement (indexed 2):

$$\begin{aligned} \hat{f}_2(W + \Delta W_1) - \hat{f}_2(W) &= (W + \Delta W_1)^T \mathbf{g}(\dot{\mathbf{x}}_2^*) - W^T \mathbf{g}(\dot{\mathbf{x}}_2^*) \\ &= -\eta \mathbf{g}^T(\dot{\mathbf{x}}_1^*) \mathbf{g}(\dot{\mathbf{x}}_2^*) \tilde{\mathbf{f}}_1 \end{aligned} \quad (3)$$

The change in the internal model output depends on experienced error and the mutual projection between evaluations of the primitives, but does not depend on the weight matrix. As the primitives depend on desired velocity, when the two movements have the same desired trajectory (for example, toward the same target), the change should be proportional to the error experienced. When the two movements are toward different targets, the change will also depend upon the breadth of the receptive fields of the primitives.

We first tested whether an error experienced in a given movement causes a proportional change in the internal model for the next movement to the same target. We asked subjects to make reaching movements while holding a manipulandum¹³ which produced viscous forces $\mathbf{f} = B\dot{\mathbf{x}}$, where $B = \{0, 13; -13, 0\} \text{ N s m}^{-1}$. Catch trials, movements during which $\mathbf{f} = 0$, were randomly interspersed among the targets. Our proxy for error was hand displacement perpendicular to target direction (perpendicular displacement; p.d.) measured 250 ms into the movement. The first movement in the field (1st in Fig. 1a) had significant error (p.d. = 2.38 cm), but with training (ct-1 in Fig. 1a) became less disturbed (p.d. = 0.45 cm). In the next movement towards this direction (90°), a catch trial (ct), there was a large error in the direction opposite the initial error, suggesting formation of an internal model¹³. In the subsequent movement to 90° (ct+1), during which the force field was present, the p.d. was substantially greater (p.d. = 1.22 cm) than in ct-1, indicating partial unlearning of the internal model as predicted by equation (2). In agreement with equation (3), there

was a significant correlation between magnitude of movement errors in the catch trial and the unlearning observed in ct+1 (Fig. 1b, $r = 0.65$). The physiological correlate of this unlearning was evident in the spatial tuning of movement-initiating muscle activations. The computational construct of an internal model predicted that spatial tuning of the electromyographic (EMG) activity of arm muscles would undergo a specific rotation with training⁵. During training, the preferred direction of this tuning gradually rotated. However, between ct-1 and ct+1 the preferred direction rotated back toward the initial orientation (Fig. 1c), indicating unlearning of the internal model. This unlearning was washed out by movement ct+3 (Fig. 1d).

We next investigated the shape of primitives underlying internal model formation by quantifying, independent of the model in equation (3), the temporal dynamics of movement errors first within and then across directions. In a sequence of random target directions, the time series of movement errors for a given direction was fitted to the following system of equations:

$$\begin{aligned} z_{n+1} &= az_n + bu_n \\ y_n &= z_n + du_n \end{aligned} \quad (4)$$

Here y represented error in the internal model as quantified by p.d., n was movement number and u indicated whether the force field

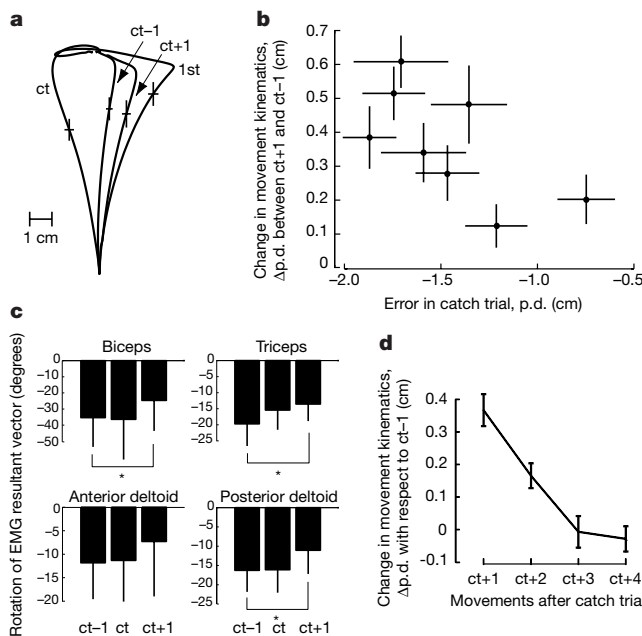


Figure 1 Catch trials induced short-term unlearning. **a**, Hand trajectories of single movements, averaged across all 40 subjects, during the first (labelled 1st), 38th (ct-1), 39th (ct) and 40th (ct+1) movements toward 90° (0° is at 3:00). Intersection of error bars indicates parallel and perpendicular displacement (p.d.) 250 ms into the movement. **b**, Jumps in p.d. between before (ct-1) and after (ct+1) catch trials versus p.d. in catch trials (ct), averaged within target directions across subjects. **c**, Orientation of preferred direction of movement-initiating EMG in ct-1, ct and ct+1. Asterisks indicate significant ($P < 0.05$) rotation of EMG preferred direction between ct-1 and ct+1. **d**, Errors in force field movements following ct. In all figures, error bars indicate 95% confidence intervals of the mean (CIM) across subjects.




Figure 2 Sensitivity to movement error across target directions. **a**, Errors in subjects' reaches (circles) (p.d. at 250 ms) toward 90°, averaged across subjects ($n = 40$). Catch trials are the large negative spikes (one is labelled ct). Lines are best fits of scalar (black) and vector (red) state-space models to the data (y in equation (4)). **b**, Average sensitivity of the internal model to errors experienced in previous movements (b in equation (4)) as a function of angular distance ϕ . Error bars (95% CIM) calculated through bootstrapping. **c**, Average generalization function $\mathbf{b}(\phi)$ for simulated adaptive controllers that constructed an internal model with Gaussians of width σ . $\mathbf{b}(\phi)$ in simulations and subjects were averaged across target directions.

was present ($u = -1$) or turned off ($u = 1$). The hidden state of the system, z , represented the amount of movement error generated by the internal model; actual error (y) also depended on whether the force field was applied. The implicit assumption in this initial model was that errors experienced in one target direction did not affect the internal model for generating movements toward other targets. The best-fit model correlated to actual errors reasonably well (Fig. 2a, black line; across directions, mean $r = 0.60$). The fit mimicked subjects' recovery from initial error, their large error in the catch trial, and their jump in error from $ct-1$ to $ct+1$. Whereas this initial model smoothly decayed after jumps in error, subjects often generated a non-monotonic change in error between catch trials. We hypothesized that this was because errors experienced in one target direction changed the internal model for other directions.

To investigate whether locally experienced errors affected other directions of movement, we expanded both u and b in equation (4) to eight-dimensional vectors. Each element of the input vector \mathbf{u} flagged recently experienced dynamics in a particular target

direction i : whether, since the last movement in the modelled target direction, a force-field movement ($\mathbf{u}(i) = -1$) or a catch trial ($\mathbf{u}(i) = 1$) had been most recently experienced, or if no movement had occurred in direction i ($\mathbf{u}(i) = 0$). Each element of \mathbf{b} , denoted $\mathbf{b}(i)$, quantified the sensitivity to errors experienced in direction i . The expanded model now accounted for subtle changes in actual movement error (Fig. 2a, red line; mean $r = 0.81$). Confidence intervals on \mathbf{b} suggested that there was a significant, nonzero influence of local errors on subsequent control in other directions. To calculate sensitivity across target directions, the elements of \mathbf{b} were re-indexed by the angular distance between the direction in which errors were experienced and the modelled movement direction. This angular distance was represented by ϕ . Averaging $\mathbf{b}(\phi)$ across movement directions (Fig. 2b) demonstrated that errors experienced in a given movement maximally influenced the internal model for that direction. This influence decayed in neighbouring directions. Surprisingly, sensitivity became significantly negative when angular distances were larger than 90° . This




Figure 3 Movement characteristics of systems that learn an internal model with velocity encoding gaussians. **a**, Simulated adaptive controllers trained in a target set without catch trials. The 11th (left) and 21st (right) movements toward 90° are shown for each controller. Learning with wide gaussians produced S-shaped movements. **b**, Adaptive controllers' estimation of the force field after training for 100 movements in a target set without catch trials. Peak velocity of typical movements is 0.35 m s^{-1} . **c**, Time course of error (p.d. at 200 ms) for simulated controllers with various gaussian widths, smoothed across a 13-movement window. Overcompensation (generation of S-shaped hand trajectories) occurs when p.d. becomes negative. **d**, Error (p.d.) averaged over the subset

of movements 65–128 during which the force field was on. With narrow gaussians, overcompensation never occurred, regardless of catch trial probability. With wide gaussians, movements became S-shaped below critical probabilities (near 17%). **e, f**, Parallel and perpendicular displacements averaged across movements and subjects. **e**, Subjects ($n = 24$) trained in target sets without catch trials, resulting in S-shaped movements. **f**, Subjects ($n = 40$) trained in target sets with 17% catch trial probability and did not produce S-shaped movements. **g**, A controller relying upon wide gaussians ($\sigma = 0.12 \text{ m s}^{-1}$) trained in target sets with 17% catch trial probability did not produce S-shaped trajectories.

indicated that when two force-field movements were separated by angular distance ϕ , if ϕ was small, then errors experienced in the first movement improved the internal model for the second movement. If ϕ was large, then errors in the first movement destructively interfered with the internal model used to generate the second movement.

To explain this result, we note that sensitivity of the internal model to experienced errors, $\mathbf{b}(\phi)$, was quantified in terms of the p.d. Both the output and the error signal of the internal model, however, are in terms of force (equations (1) and (2)). Because the force field is linear in velocity, the direction of force error corresponding to positive (clockwise) p.d. towards one target opposes the direction of force error corresponding to a positive p.d. towards the opposite target. Interpreting the sensitivity of subjects' internal model (Fig. 2b) through the adaptation rule (equation (3)) suggests that both the positive values of \mathbf{b} for $-45^\circ < \phi < 45^\circ$ and the negative values of \mathbf{b} for large ϕ correspond to the same direction of force compensation. From this we deduced that the mutual projection $\mathbf{g}^T(\dot{\mathbf{x}}_i^*)\mathbf{g}^T(\dot{\mathbf{x}}_{i+1}^*)$ declines but always remains positive as the angular distance between two movements increases. This result rules out bases that encode velocity space linearly. Furthermore, because information experienced in each direction most strongly affects that direction and its neighbour less so, basis functions that have specific regions of preferred activity are more likely to underlie learning than global representations of dynamics.

We therefore investigated what conditions on $\mathbf{g}(\dot{\mathbf{x}}^*)$ were sufficient to generate the generalization function $\mathbf{b}(\phi)$. A salient property of cells in the motor system is their directional tuning¹⁴ and modulation with hand speed¹⁵. In the cerebellum, a region which

lesion^{16–19} and functional imaging studies^{20,21} have linked to learning and control of arm dynamics, many Purkinje cells simultaneously encode the direction and speed components of velocity²². These cells broadly encode hand velocity during planar reaching, firing maximally at preferred velocities distributed in velocity space. This encoding precedes in time the actual movement, suggesting that these cells encode desired velocity. The behaviour of each cell k could therefore be represented as a gaussian with a centre located at position \mathbf{c}_k in desired velocity space. We simulated a controller attached to a biomechanical model of the arm that learned an internal model with basis functions:

$$g_k(\dot{\mathbf{x}}^*) = \exp(-|\dot{\mathbf{x}}^* - \mathbf{c}_k|^2/2\sigma^2) \quad (5)$$

where σ is the standard deviation of the gaussian. To accommodate the possibility that the exact shape of $\mathbf{b}(\phi)$ depended on the training paradigm, we trained subjects and the simulated controller with the identical set of targets and catch trials. A crucial component of the simulations was σ , the width of the primitives. When the gaussians were narrow, the time series of errors generated by the simulation showed a spike after each catch trial and a smooth decay afterwards, similar to the scalar-input state–space model fit (Fig. 2a, black line) but unlike the performance of our subjects. Simulations driven by broad gaussians, however, produced non-monotonic changes that mimicked subjects' actual patterns of adaptation. Simulation results were fitted with equation (4) to produce the generalization function $\mathbf{b}(\phi)$ (Fig. 2c). The $\mathbf{b}(\phi)$ generated with narrow gaussians rapidly dropped to zero as ϕ changed from zero. Learning with wide gaussians, however, showed a generalization that was very similar to actual subject performance, including negative sensitivity for large




Figure 4 Learning with wide gaussians imposes limits on adaptation. **a**, Correlation between actual force field and internal model acquired after 50 targets by adaptive controllers with various width gaussians, versus the frequency of the nonlinear force field (m in equation (6)). **b**, A high spatial frequency field ($m = 4$) in which subjects and a simulated controller trained. **c**, The internal model learned by the controller with $\sigma = 0.12 \text{ m s}^{-1}$ gaussians after training for 50 movements. Note that for directions 22.5° and 157.5°, the controller predicts the field to be resistive, whereas actual forces are

assistive. **d**, Hand velocities parallel to target direction of three subjects trained in the same high-spatial-frequency field in the same target set as the simulated controller. Plotted are velocities in baseline null field movements (mean \pm standard deviation) and during catch trials in targets 51 and 58 toward 22.5° and 157.5°. In catch trials, the hand is moving significantly faster, consistent with an internal model that expects a resistive field.

ϕ . The correlation between $\mathbf{b}(\phi)$ in the simulations and the subject data was strongest for $\sigma = 0.12 \text{ m s}^{-1}$.

We next used the model to predict behaviour beyond the data set with which the primitives were estimated. We noted that gaussian width influences how force estimation generalizes across both directions and speeds. Simulations predicted that when learning relied upon wide gaussians, reaching movements would not monotonically converge onto a straight line desired trajectory but would become S-shaped (Fig. 3a). Whereas the force field was linear in velocity, wide gaussians produced an approximation that overestimated forces at low speeds and underestimated forces at high speeds (Fig. 3b). Overestimation of the forces resulted in overcompensation of the field early in the movement; the magnitude of overcompensation depended on the gaussian width (Fig. 3c). With narrow gaussians, the simulated internal model did not overcompensate, but with wide gaussians movements became S-shaped. To test this prediction, we trained 24 subjects in target sets without catch trials. Movements of subjects were S-shaped (Fig. 3e), similar to movements made by simulations that learned with wide gaussians (Fig. 3a).

Simulations further predicted that the probability of catch trials influenced whether movements would become S-shaped (Fig. 3d). If catch trials occurred with 17% probability, then even with gaussians of $\sigma = 0.12 \text{ m s}^{-1}$ there should be sufficient unlearning caused by each catch trial such that hand trajectories would converge toward a straight line, without overcompensation (Fig. 3g). We tested this prediction by training subjects in target sets with 17% catch trial probability. As predicted, subjects did not show overcompensation (Fig. 3f).

Because approximation of a high-frequency signal with low-frequency bases generally results in poor representation, we next explored the limitations of a system that learns with wide gaussians. We simulated learning of nonlinear force fields:

$$\mathbf{f} = -13\sqrt{\dot{x}^2 + \dot{y}^2} \begin{bmatrix} -\sin(m\theta) \\ \cos(m\theta) \end{bmatrix} \quad (6)$$

$$\theta = \arctan(\dot{y}/\dot{x})$$

where \dot{x} and \dot{y} were components of hand velocity in Cartesian space. When $m = 1$, the field was the curl pattern¹³ learned by subjects above. As m increased, the field's spatial frequency increased. For various σ , we simulated movements to 50 targets and correlated the internal model learned by the simulation to the actual force field (Fig. 4a). As the field's spatial frequency increased, the accuracy of the internal model decreased for all basis function widths. However, the learning capability of wider bases collapsed at lower frequencies. This agrees with the recent finding that humans demonstrated a lesser ability to adapt in higher-spatial-frequency force fields²³.

To illustrate this deficiency, we trained an adaptive controller ($\sigma = 0.12 \text{ m s}^{-1}$) for 50 movements in a high-spatial-frequency field ($m = 4$, Fig. 4b). Because approximation is performed with wide bases, the internal model learned by the controller (Fig. 4c) cannot represent faithfully the rapidly changing forces. In particular, the simulation predicts that in movements toward 22.5° and 157.5° the internal model expects resistive forces where the force field is assistive. We tested for this prediction in three subjects by training each with the same random pattern of targets presented to the simulation, then presenting two catch trials toward 22.5° and 157.5°. Subjects behaved as though they were expecting a resistive field in those directions, as illustrated by their hand velocities in the catch trials (Fig. 4d).

Errors in learning dynamics of arm movements suggest that the brain composes motor commands with computational elements that are broadly tuned to arm velocity. When expressed in polar coordinates, the gaussians exhibit a preferred direction of movement, much like the cosine tuning curves typically associated with

cells in the motor system. However, our inferred bases have on average a half-width at half-height value of about 40°, significantly less than the 90° value required of cosines. Recent results²⁴ have demonstrated that tuning curves in monkey motor cortex have a median width of 56°, also much narrower than cosines. Motor cortical cells, however, have been reported to encode hand speed linearly¹⁵. Learning a linear force field with bases that have cosine directional tuning and linear speed tuning results in an internal model that does not produce the S-shaped movements observed in our subjects. Wide gaussians predict this unusual behaviour. They also explain why humans generalize sublinearly to fast movements after training in a linear force field at slow speeds⁸. The nonlinear encoding of speed inherent to gaussians resemble tuning properties of Purkinje cells that encode arm movements in the cerebellum²². Although several investigators have proposed a major role for the cerebellum in learning of internal models^{25–27}, our results suggest a link between patterns of generalization and firing properties of cells in this area. Since the output of the cerebellum partially affects the motor cortex²⁹, the finding that the preferred direction of motor cortical cells rotates during learning of force fields²⁸ may be a consequence of changing input from the cerebellum³⁰. □

Methods

Three groups of right-handed normal human subjects were trained to make movements while holding the handle of a lightweight robotic arm. All movements were toward a pseudorandomly chosen target, then back to a centre target. Targets appeared at 0°, 45°, ... 315°, at 10-cm displacement. Desired movement duration was $500 \pm 50 \text{ ms}$. Timing feedback was provided by changing the target colour. All subjects initially practised the task without any perturbing force (the null field). The first group of subjects ($n = 40$) trained in target sets of 192 movements during which the force field $B = [0, 13; -13, 0] \text{ N s m}^{-1}$ was applied in 83% of the targets. The remaining 17% of targets (pseudorandomly selected) were catch trials during which the force field was turned off. In 13 of these subjects, EMG from anterior and posterior deltoid, biceps and triceps were measured with surface electrodes, amplified, filtered, r.m.s. (root mean square)-rectified, averaged from 50 to 100 ms into the movement, multiplied by unit vectors pointing toward movement direction, and summed across movement directions to produce a preferred direction for each muscle⁵. The second group of subjects ($n = 24$) trained in the same target sets as above, but with the force field always on (no catch trials). A final group of subjects ($n = 3$) trained for 58 movements in a higher-frequency force field (equation (6); $m = 4$), receiving catch trials only on targets 51 and 58 in directions 22.5° and 157.5°.

A simulated anthropomorphic controller¹³ made movements to the same sequence of targets experienced by subjects. The controller's internal model learned to map desired hand velocities into forces. Desired hand trajectories were minimum jerk, 0.5 s in duration, and 10 cm long. The internal model approximated the imposed force field (equation (1)) with gaussian bases (equation (5)); controllers were differentiated by σ . Gaussian centres spanned a range of desired velocities (-0.5 to 0.5 m s^{-1} in the x - and y -direction) spaced one σ apart. Sensitivity analyses suggested there were no significant changes in the results when the density of the bases was increased by an order of magnitude. Weights were initially randomized, then updated using equation (2) with $\eta = 0.0025$. To compare the output of gaussian basis functions with cosine tuning curves, the unweighted output of each gaussian was averaged during movement time. The collection of averaged outputs across movement directions formed a tuning curve for each gaussian primitive. Tuning curves were aligned to the preferred direction of each gaussian and averaged across all gaussians, from which the half-width at half-height value was calculated.

Received 17 May; accepted 5 July 2000.

- Mussa-Ivaldi, F. A., Giszter, S. F. & Bizzi, E. Linear combinations of primitives in vertebrate motor control. *Proc. Natl Acad. Sci. USA* **91**, 7534–7538 (1994).
- Schaal, S. & Atkeson, C. G. Constructive incremental learning from only local information. *Neural Comput.* **10**, 2047–2084 (1998).
- Shadmehr, R. & Mussa-Ivaldi, F. A. Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.* **14**, 3208–3224 (1994).
- Ghez, C., Krakauer, J. W., Sainburg, R. L. & Ghilardi, M. F. in *The New Cognitive Neurosciences* (ed. Gazzaniga, M. S.) 501–514 (MIT Press, Cambridge, Massachusetts, 2000).
- Thoroughman, K. A. & Shadmehr, R. Electromyographic correlates of learning internal models of reaching movements. *J. Neurosci.* **19**, 8573–8588 (1999).
- Gandolfo, F., Mussa-Ivaldi, F. A. & Bizzi, E. Motor learning by field approximation. *Proc. Natl Acad. Sci. USA* **93**, 3843–3846 (1996).
- Condit, M. A., Gandolfo, F. & Mussa-Ivaldi, F. A. The motor system does not learn the dynamics of the arm by rote memorization of past experience. *J. Neurophysiol.* **78**, 554–560 (1997).
- Goodbody, S. J. & Wolpert, D. M. Temporal and amplitude generalization in motor learning. *J. Neurophysiol.* **79**, 1825–1838 (1998).
- Condit, M. A. & Mussa-Ivaldi, F. A. Central representation of time during motor learning. *Proc. Natl Acad. Sci. USA* **96**, 11625–11630 (1999).
- Kawato, M. Adaptation and learning in control of voluntary movement by the central nervous system. *Adv. Robotics* **3**, 229–249 (1989).

11. Sanner, R. M. & Koshi, M. A mathematical model of the adaptive control of human arm motions. *Biol. Cybern.* **80**, 369–382 (1999).
12. Bhushan, N. & Shadmehr, R. Computational architecture of human adaptive control during learning of reaching movements in force fields. *Biol. Cybern.* **81**, 39–60 (1999).
13. Shadmehr, R. & Brashers-Krug, T. Functional stages in the formation of human long-term motor memory. *J. Neurosci.* **17**, 409–419 (1997).
14. Georgopoulos, A. P., Schwartz, A. B. & Kettner, R. E. Neural population coding of movement direction. *Science* **233**, 1416–1419 (1986).
15. Moran, D. W. & Schwartz, A. B. Motor cortical representation of speed and direction during reaching. *J. Neurophys.* **82**, 2676–2692 (1999).
16. Conrad, B., Matsunami, K., Meyer-Lohmann, J., Wiesendanger, M. & Brooks, V. B. Cortical load compensation during voluntary elbow movements. *Brain Res.* **71**, 507–514 (1974).
17. Goodkin, H. P., Keating, J. G., Martin, T. A. & Thach, W. T. Preserved simple and impaired compound movement after infarction in the territory of the superior cerebellar artery. *Can. J. Neurol. Sci.* **20**, S93–S104 (1993).
18. Bastian, A. J., Martin, T. A., Keating, J. G. & Thach, W. T. Cerebellar ataxia: abnormal control of interaction torques across multiple joints. *J. Neurophysiol.* **76**, 492–509 (1996).
19. Lang, C. E. & Bastian, A. J. Cerebellar subjects show impaired adaptation of anticipatory EMG during catching. *J. Neurophysiol.* **82**, 2108–2119 (1999).
20. Nezafat, R., Shadmehr, R. & Holcomb, H. H. Neural correlates of motor memory retention over a 4 week period as quantified through PET. *Soc. Neurosci. Abs.* **25**, 2178 (1999).
21. Imamizu, H. *et al.* Human cerebellar activity reflecting an acquired internal model of a new tool. *Nature* **403**, 192–195 (2000).
22. Coltz, J. D., Johnson, M. T. V. & Ebner, T. J. Cerebellar Purkinje cell simple spike discharge encodes movement velocity in primates during visuomotor arm tracking. *J. Neurosci.* **19**, 1782–1803 (1999).
23. Matsouka, Y. Models of generalization in motor control. PhD thesis, MIT (1998).
24. Amirkian, B. & Georgopoulos, A. P. Directional tuning profiles of motor cortical cells. *Neurosci. Res.* **36**, 73–79 (2000).
25. Houk, J. C. & Wise, S. P. Distributed modular architectures linking basal ganglia, cerebellum, and cerebral cortex: their role in planning and controlling action. *Cereb. Cortex* **5**, 95–110 (1995).
26. Wolpert, D. M. & Kawato, M. Multiple paired forward and inverse models for motor control. *Neural Networks* **11**, 1317–1329 (1998).
27. Spoelstra, J., Schweighofer, N. & Arbib, M. A. Cerebellar learning of accurate predictive control for faster-reaching movements. *Biol. Cybern.* **82**, 321–333 (2000).
28. Gandolfo, F., Li, C. S. R., Benda, B. J., Schioppa, C. P. & Bizzi, E. Cortical correlates of learning in monkeys adapting to a new dynamical environment. *Proc. Natl Acad. Sci. USA* **97**, 2259–2263 (2000).
29. Holdefer, R. N., Miller, L. E. & Houk, J. C. Functional connectivity between cerebellum and primary motor cortex in the awake monkey. *J. Neurophysiol.* **84**, 585–590 (2000).
30. Martin, J. H., Cooper, S. E., Hacking, A. & Ghez, C. Differential effects of deep cerebellar nuclei inactivation on reaching and adaptive control. *J. Neurophysiol.* **84**, 1886–1899 (2000).

Acknowledgements

We thank M. A. Smith, O. Donchin, and R. Nezafat. The work was supported by grants from the Office of Naval Research and the National Institutes of Health (to R.S.), and a research traineeship from the National Science Foundation (K.A.T.).

Correspondence should be addressed to R.S. (e-mail: reza@bme.jhu.edu).

Netrin-1-mediated axon outgrowth and cAMP production requires interaction with adenosine A2b receptor

Véronique Corset*, Kim Tuyen Nguyen-Ba-Charvet†, Christelle Forcet*, Emmanuel Moyses*, Alain Chédotal† & Patrick Mehlen*

* Laboratoire Apoptose et Différenciation, label 'La Ligue' – Centre de Génétique Moléculaire et Cellulaire, CNRS UMR 5534, Université Lyon 1, 69622 Villeurbanne, France

† INSERM U106, Hôpital de la salpêtrière, 75651 Paris cedex 13, France

The netrins, a family of laminin-related secreted proteins, are critical in controlling axon elongation and pathfinding^{1–4}. The DCC (for deleted in colorectal cancer) protein was proposed as a receptor for netrin-1 in the light of many observations including the inhibition of netrin-1-mediated axon outgrowth and attraction in the presence of an anti-DCC antiserum^{5–7}, the similitude of nervous system defects in DCC and netrin-1 knockout mice^{4,8} and the results of receptor swapping experiments⁹. Previous studies

have failed to show a direct interaction of DCC with netrin-1 (ref. 10), suggesting the possibility of an additional receptor or co-receptor. Here we show that DCC interacts with the membrane-associated adenosine A2b receptor, a G-protein-coupled receptor that induces cAMP accumulation on binding adenosine¹¹. We show that A2b is actually a netrin-1 receptor and induces cAMP accumulation on binding netrin-1. Finally, we show that netrin-1-dependent outgrowth of dorsal spinal cord axons directly involves A2b. Together our results indicate that the growth-promoting function of netrin-1 may require a receptor complex containing DCC and A2b.

To identify direct partners of DCC, we carried out a two-hybrid screen of a human brain complementary DNA library using the intracellular domain of DCC as bait. The screen revealed the putative binding of different proteins to the intracellular domain of DCC. Some of these partners, like caspase-3, have been described elsewhere^{12,13}. We also observed putative binding of actin and dynactin to DCC (data not shown); this could be important because actin and dynactin are cytoskeleton proteins that are involved in neurite elongation and guidance. But here we focused our attention on the relevance of the DCC binding to a 23-residue fragment corresponding to the last 23 amino acids (310–333) of the intracellular domain of the adenosine A2b receptor (Fig. 1a). A2b is a member of the family of the adenosine-specific receptors¹¹, a family which was originally divided into four subtypes—the A1, A2a, A2b and A3 receptors—on the basis of their ability to inhibit (through A1 or A3) or stimulate (through A2a or A2b) adenylate cyclase activity in response to adenosine binding¹⁴. The A2b receptor couples to either G_q or G_s, leading to a phospholipase C-dependent pathway¹⁵ or an accumulation of cAMP¹⁶, respectively. The low affinity of the A2b receptor for adenosine as compared with A2a questioned the role of this receptor *in vivo*.

We confirmed the DCC/A2b interaction identified in the two-hybrid screen by co-immunoprecipitation (Fig. 1b). In 293T cells co-transfected with A2b and DCC, a DCC/A2b interaction was barely detectable in the absence of netrin-1, whereas there was a very strong binding when netrin-1 was present (Fig. 1b). Hence, DCC




Figure 1 DCC interacts with the adenosine A2b receptor. **a**, Two-hybrid screen of DCC-IC revealed interaction with A2b. Transformed Y190 yeast growth in the selection medium is shown. **b–d**, Co-immunoprecipitations were performed on 293T cells co-transfected with combinations of the A2b encoding plasmid Prc/CMV-A2b-Flag (A2b), the full-length DCC encoding plasmid pDCC-CMV-S (DCC) and the netrin-1 encoding plasmid pGNET1myc (net). The pull-down assays were done with anti-Flag (**b, d**) or anti-c-Myc (**c**) antibody. In **d**, DCC-IC fused to a myristoylation signal¹² was expressed instead of DCC full-length. 'tot.' and 'IP' indicate DCC immunoblots before and after the pull-down assays, respectively.

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