

between the birds was their experience in the earlier observational learning study¹⁴—observer, pilferer or observer + pilferer. These categories are directly related to the specific experiences in the previous experiment, not the developmental histories of the birds. Although we cannot state that birds in the pilferer group had never observed other birds cache they did not do so in the previous study or in the current one. Importantly, the observer birds have never pilfered another bird's caches.

The observer + pilferer ($n = 7$) group was first tested at the University of California, Davis (19–26 July 2000), with a replicate study performed at the University of Cambridge (15 January to 16 February 2001). Birds in the observer ($n = 7$) and pilferer ($n = 7$) groups were tested at Cambridge (15 January to 16 February 2001). The interleaved trials study (observer + pilferer group only) was performed at Cambridge (27–28 February 2001). All of the birds were housed individually in cages (45 × 76 × 76 cm). In Davis, birds were housed in an outdoor aviary. In addition to natural lighting, we provided fluorescent strip lighting. In Cambridge, the birds were housed in the same cages as Davis, but these were placed in a quarantined indoor room in which only fluorescent lighting was provided. Birds were fed a mixture of Iams mini-chunk dog food biscuits and peanuts, both of which were provided in powdered form to ensure that the birds could not cache outside of the experiment. Water was provided *ad libitum*.

Apparatus

The fronts of the individual cages were made of aluminium wire and the sides were made of solid aluminium. The storer was placed in a cage located adjacent to a second cage containing the observer, with a 20-cm-wide gap between the two cages. On the back of each cage was a 25 cm² perspex panel, which allowed the storer and observer to see one another. Caching trays^{10–13} were constructed from sand-filled plastic ice-cube trays, each containing a 2 × 8 array of moulds that were attached to a wooden board. Each caching tray was made unique by attaching different configurations of Lego (Netfield) bricks onto the board behind the ice-cube tray.

Procedure

Birds were deprived of food overnight. Caching trials started at 10:00 (Davis) or 11:00 (Cambridge) the following day. During every 15-min caching trial ($n = 3$ per caching treatment, observer + pilferer group in Davis and observer, pilferer and observer + pilferer groups in Cambridge; $n = 2$ pairs, observer + pilferer group interleaved trials in Cambridge) each subject received one sand-filled caching tray and a bowl containing 50 wax worms. During the unobserved caching treatment (in private), a towel covered the back of the cage so that the view of the observer bird was completely obscured. During the observed caching treatment, the observer had a clear view of the storer. At the end of each caching trial, the tray and food bowl were removed from each cage. The experimenters recorded the number and location of wax worms cached in the tray. Any extraneous food that the birds did not cache was removed. The number of worms cached therefore refers to the number of worms that remained in the tray at the end of the caching trial. Observers were given 15 min in which they were allowed to eat their maintenance diet at the end of the caching trial, and each bird was given three wax worms. In Davis and Cambridge, two caching trays were placed inside the subject's cage during recovery trials. One of the trays was unfamiliar to the bird. The other tray contained the previously cached worms and was placed in its original location. Three caching trays were placed inside the subject's cage during recovery trials for the interleaved trials study. One tray was unfamiliar to the bird, whereas the other two trays contained the previously cached worms, which were placed in the same locations as they were placed during the observed tray and in private tray caching treatments.

The subjects were allowed to recover the cached food items for 10 min. The number and location of the searches was recorded by direct observation, as well as the number of caches recovered and whether these were eaten or re-cached. We also noted the location of the re-cached food items, and whether they were made in new sites compared to those places the bird had cached in during the previous caching trial (old site). We calculated the number of searches to find the first cache (recovery accuracy), the proportion of worms recovered and the proportion of those recoveries that were re-cached. All 21 birds had previous experience with the presence of towels on their cages for at least three trials each. They were restricted from caching during these trials, but were given powdered food and three wax worms.

Analysis

We used the Wilcoxon matched-pairs test on the mean values summed across trials for each caching treatment. As the number of items recovered depends on the number of items cached, and the number of items re-cached depends on the number of caches that are recovered, we also analysed the proportion of caches that were recovered and the proportion of those re-cached. For all analyses, alpha was set at 0.05.

Received 23 July; accepted 20 September 2001.

- Zentall, T. R. & Galef, B. G. (eds) *Social Learning. Psychological and Biological Perspectives* (Lawrence Erlbaum Associates, London, 1998).
- Shettleworth, S. J. in *Behavioral Brain Research in Naturalistic and Semi-naturalistic Settings* (eds Alleva, E., Fasolo, A., Lipp, H.-P. & Nadel, L.) 158–179 (Kluwer Academic, The Hague, 1995).
- Vander Wall, S. B. *Food Hoarding in Animals* (Univ. Chicago Press, Chicago, 1990).
- Clarkson, K., Eden, S. F., Sutherland, W. J. & Houston, A. I. Density dependence and magpie food hoarding. *J. Anim. Ecol.* **55**, 111–121 (1986).
- Gibb, J. A. Populations of tits and goldcrests and their food supply in pine populations. *Ibis* **102**, 163–208 (1960).
- Bednekoff, P. A. & Balda, R. P. Observational spatial memory in Clark's nutcrackers and Mexican jays. *Anim. Behav.* **52**, 833–839 (1996).

- Heinrich, B. & Pepper, J. W. Influence of competitors on caching behavior in the common raven, *Corvus corax*. *Anim. Behav.* **56**, 1083–1090 (1998).
- Bugnyar, T. & Kotrschal, K. Do ravens manipulate the others' attention in order to prevent or achieve social learning opportunities? *Adv. Ethol.* **36**, 106 (2001).
- Heinrich, B. *Mind of the Raven* (Harper Collins, New York, 1999).
- Clayton, N. S. & Dickinson, A. D. Episodic-like memory during cache recovery by scrub jays. *Nature* **395**, 272–278 (1998).
- Clayton, N. S. & Dickinson, A. D. Scrub jays (*Aphelocoma coerulescens*) remember the relative time of caching as well as the location and content of their caches. *J. Comp. Psychol.* **113**, 403–416 (1999).
- Clayton, N. S. & Dickinson, A. D. Memory for the contents of caches by Scrub Jays. *J. Exp. Psychol. Anim. Behav. Proc.* **25**, 82–91 (1999).
- Clayton, N. S., Yu, K. & Dickinson, A. D. Scrub jays (*Aphelocoma coerulescens*) can form integrated memory for multiple features of caching episodes. *J. Exp. Psychol. Anim. Behav. Proc.* **27**, 17–29 (2001).
- Clayton, N. S., Griffiths, D. P., Emery, N. J. & Dickinson, A. D. Episodic-like memory in animals. *Phil. Trans. R. Soc. Lond. B* **356**, 1483–1491 (2001).
- Goodwin, D. Further observations on the behaviour of the jay. *Ibis* **98**, 186–219 (1956).
- Stevens, T. A. & Krebs, J. R. Retrieval of stored seeds by marsh tits (*Parus palustris*) in the field. *Ibis* **128**, 513–515 (1984).
- Hampton, R. R. & Sherry, D. F. The effects of cache loss on choice of cache sites in the black-capped chickadee. *Behav. Ecol.* **5**, 44–50 (1994).
- Baker, M. C. & Anderson, P. Once-pilfered cache sites not avoided by black-capped chickadees. *Anim. Behav.* **49**, 1599–1602 (1995).
- Hare, B., Call, J., Agnetta, B. & Tomasello, M. Chimpanzees know what conspecifics do and do not see. *Anim. Behav.* **59**, 771–785 (2000).
- Hare, B., Call, J. & Tomasello, M. Do chimpanzees know what conspecifics know? *Anim. Behav.* **61**, 139–151 (2001).
- Whiten, A. & Byrne, R. W. Tactical deception in primates. *Behav. Brain Sci.* **11**, 233–244 (1988).
- Suddendorf, T. & Corballis, M. C. Mental time travel and the evolution of the human mind. *Genet. Soc. Gen. Psychol. Monogr.* **123**, 133–167 (1997).
- Heyes, C. M. Theory of mind in nonhuman primates. *Behav. Brain Sci.* **21**, 101–148 (1998).
- Clayton, N. S. & Dickinson, A. D. Motivational control of caching behaviour in the scrub jay, *Aphelocoma coerulescens*. *Anim. Behav.* **57**, 435–444 (1999).

Acknowledgements

This work was supported by the National Institutes of Health, National Institute on Aging and Whitehall foundation grants to N.S.C. N.J.E was funded by a Medical Research Council programme grant to E. B. Keverne. We thank S. Hettige, D. Jennings and V. R. Metcalf for help in running the experiments. We also thank the University of California, Davis, for allowing us to conduct the first trials there and for providing the necessary facilities. We thank S. Baron-Cohen, E. B. Keverne, K. N. Laland, D. I. Perrett and C. J. Saldanha for comments on the manuscript. We are especially grateful to T. J. Bussey, A. Dickinson and N. J. Mackintosh for discussion and comments on the manuscript.

Correspondence and requests for materials should be addressed to N.S.C. (e-mail: nsc22@cam.ac.uk).

The central nervous system stabilizes unstable dynamics by learning optimal impedance

Etienne Burdet*†§||, Rieko Osu*†||, David W. Franklin§¶, Theodore E. Milner¶ & Mitsuo Kawato*†§

* Department of Mechanical Engineering, National University of Singapore, 119260, Singapore

† Kawato Dynamic Brain Project, ERATO, JST, HIKARIDAI, Seika-cho, Soraku-gun, Kyoto, 619 0288, Japan

§ ATR Human Information Science Laboratories, HIKARIDAI, Seika-cho, Soraku-gun, Kyoto, 619 0288, Japan

¶ School of Kinesiology, Simon Fraser University, Burnaby, British Columbia V5A 1S6, Canada

|| These authors contributed equally to the work

To manipulate objects or to use tools we must compensate for any forces arising from interaction with the physical environment. Recent studies indicate that this compensation is achieved by learning an internal model of the dynamics^{1–6}, that is, a neural representation of the relation between motor command and

movement^{5,7}. In these studies interaction with the physical environment was stable, but many common tasks are intrinsically unstable^{8,9}. For example, keeping a screwdriver in the slot of a screw is unstable because excessive force parallel to the slot can cause the screwdriver to slip and because misdirected force can cause loss of contact between the screwdriver and the screw. Stability may be dependent on the control of mechanical impedance in the human arm because mechanical impedance can generate forces which resist destabilizing motion. Here we examined arm movements in an unstable dynamic environment created by a robotic interface. Our results show that humans learn to stabilize unstable dynamics using the skilful and energy-efficient strategy of selective control of impedance geometry.

How do we succeed in performing mechanically unstable tasks? One way would be to stabilize the arm by means of feedback control. However, neural feedback, whether operating through involuntary (reflex) or voluntary commands to our muscles, is delayed by the sensory feedback pathways, which would tend to increase rather than reduce instability^{10,11}. Alternatively, stability might be achieved by controlling mechanical impedance (resistance to imposed motion); specifically, the spring-like property of muscles^{12,13}. Demonstrating impedance control is not trivial, because it is necessary to show that the shape or orientation of the impedance can be voluntarily modified independently of the force applied by the hand. Modifications in impedance that depend on changes in applied force do not constitute proof of impedance control because muscle stiffness inherently scales with activation level^{14,15}. To demonstrate impedance control, we compared multi-joint arm movements in a null field (NF) and in a divergent force field (DF), which produced an unstable interaction with the arm, but required no change in applied force relative to the NF.

Subjects made horizontal point-to-point movements away from the body, along the *y* axis of our coordinate system (Fig. 1a). The hand was linked by means of a stiff brace to a handle at the end of a robotic manipulandum (PFM) that exerted computer-controlled forces during movement (see Methods). The DF produced a negative elastic force perpendicular to the target direction. The

robot produced no force when trajectories followed the *y* axis, but the hand was pushed away whenever it deviated from the *y* axis (red arrows in Fig. 1a). The force (F_x, F_y) exerted on the hand was

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} \beta x \\ 0 \end{bmatrix} \quad (1)$$

where $\beta > 0$ (N m^{-1}) was chosen to be larger than the stiffness of the arm measured in NF movements so as to produce instability.

The NF movements (green in Fig. 1b) had trajectories that were approximately straight. Almost all trials ended on target and the distribution of endpoints was bell-shaped. However, the initial movement direction varied slightly from trial to trial owing to motor output variability. The DF amplifies such variations by pushing the hand with a force proportional to the deviation from the *y* axis (red arrows in Fig. 1a). Consequently, the initial trials in the DF exhibited unstable behaviour, diverging widely either to the right or to the left. Most trials ended either on the left or right safety barrier (black vertical lines in the left plot of Fig. 1c) with a roughly symmetrical distribution of positions. With practice, however, subjects gradually became proficient at producing straighter trajectories, similar to those in the NF (compare Fig. 1b and the middle plot of Fig. 1c). Most trials then ended on target with a bell-shaped endpoint distribution. The hand-path error (see Methods, equation (2)) decreased significantly between the first five trials and the last five trials ($P = 0.002$; Fig. 2a), indicating that the subjects learned to adapt to the dynamics of the DF. At the end of the learning trials, no significant differences in the hand-path error were found between

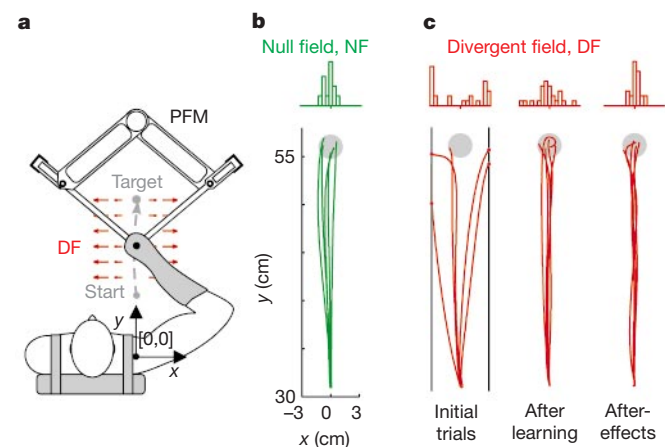


Figure 1 Experimental set-up and learning in a divergent force field (DF). **a**, The PFM exerts forces on the hand during horizontal point-to-point arm movements away from the body. The red arrows graphically depict the DF. **b**, Five consecutive trajectories of one subject are shown for movements in a null force field (NF, green). At the top of trajectory plot, a histogram shows the distribution of endpoint positions (based on curvature greater than 0.03 m^{-1}) for the corresponding five movements of all subjects. **c**, Learning in the DF (red). Five trajectories of the same subject are shown, which include the first five trials (left), the last five trials (middle), and five consecutive after-effect trials (right). For safety reasons the DF was turned off when the trajectory deviated more than 3 cm from the *y* axis. The black lines indicate this safety zone.

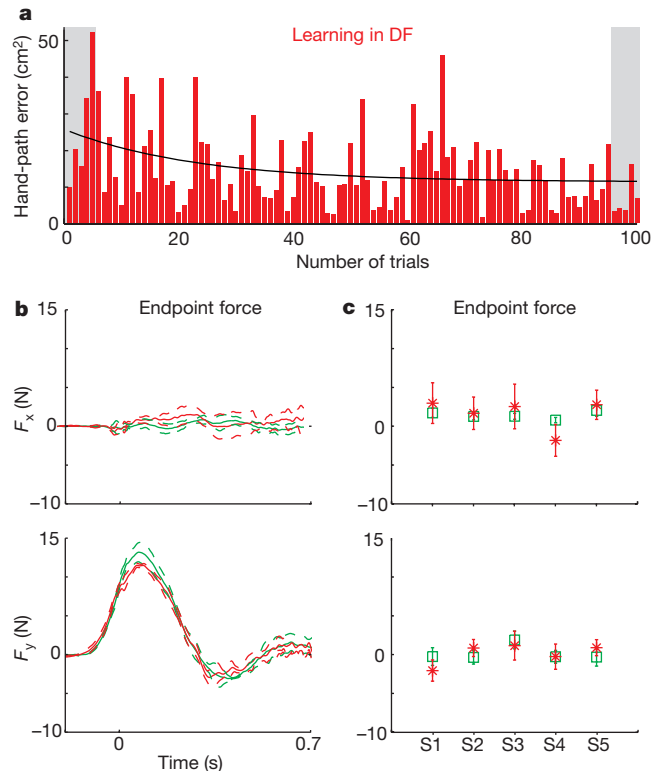


Figure 2 Error evolution during learning the DF, endpoint force–time profiles after learning for the same subject as in Fig. 1, and forces at the movement midpoint for all subjects. The hand-path error is defined in the Methods. **a**, Hand-path error as a function of the number of learning trials for the DF. The solid curve is a least-squares fit of an exponential function to the data. **b**, Mean endpoint force profiles with standard deviations in the *x* and *y* directions after learning, for NF (green) and DF (red). The force is similar in the DF and NF. **c**, The mean *x* and *y* forces and standard deviations at the midpoint of the trajectory in the NF (green) and DF (red) for all five subjects.

movements in the NF and DF ($P = 0.906$), indicating that movements had become stable and were similar under the two conditions. This is also indicated by the force applied to the PFM by the arm in the DF, which was not significantly different from that in the NF, after learning (Fig. 2b, c).

To examine the after-effects of learning the DF, the force field was unexpectedly removed on selected trials after learning. These after-effect movements were characterized by trajectories that hardly deviated from the y axis with endpoints concentrated in the middle of the target (Fig. 1c, right). The hand-path error revealed that they were even straighter than NF trajectories ($P = 0.001$).

How did subjects overcome the instability of their movements? Because of the unpredictability of the direction of the disturbance a forward or inverse dynamics model^{5,7} could not have been used. In contrast, an increase in mechanical impedance would provide increased resistance to a disturbance regardless of its direction. Perhaps subjects used generalized muscle coactivation to achieve an isotropic increase in stiffness at the hand. This has been demonstrated previously during posture¹⁶. An analogous strategy is commonly employed in single-joint tasks when the environment is destabilizing^{17,18}. On the other hand, the stiffness might have been selectively controlled in the direction of the instability. Such control of limb impedance was first proposed 16 years ago¹², but has yet to be demonstrated.

These possibilities were investigated by examining the stiffness at the hand after learning (see Methods). We measured stiffness around the midpoint of movement paths, using a technique¹⁹ that provides an unbiased and accurate estimate of stiffness and requires fewer trials than previous methods^{20,21}. On random trials, the PFM briefly displaced the hand by a constant distance from a prediction of the undisturbed trajectory¹⁹. Stiffness was estimated from the restoring force, produced in response to the displacement, divided by the amplitude of the displacement.

Stiffness at the hand in the DF was modified in shape and orientation relative to stiffness in the NF (Fig. 3, compare red and green data). Although the x component of stiffness was significantly larger in the DF than in the NF (Fig. 4a, compare red and green data), the y component of stiffness in the DF was not significantly different from that in the NF (Fig. 4b, compare red and green data).

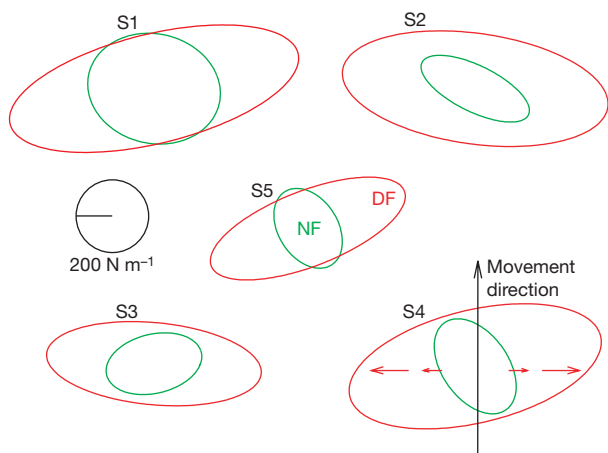


Figure 3 Adaptation of stiffness geometry to the DF. Stiffness ellipses—the graphical depiction of the elastic restoring force corresponding to a unit displacement of the hand—are shown for five subjects in the NF and DF. The long and short axes of the ellipse represent the directions of maximal and minimal stiffness, respectively. The stiffness in the DF (red) changed relative to the NF (green). We note that in the DF the stiffness increased significantly in the direction of the instability (shown by small red arrows for subject S4), but there was relatively little change along the movement direction (see also Fig. 4).

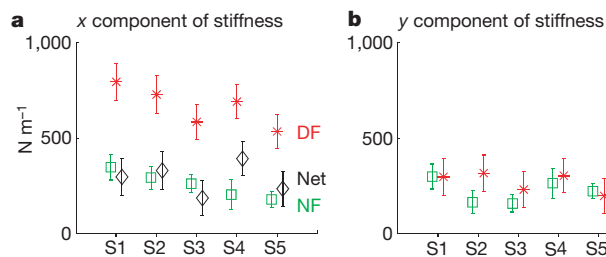


Figure 4 Stiffness and stability. **a**, The x components of hand stiffness in the NF (green), hand stiffness in the DF (red) and net stiffness of the hand and DF environment (Net, black) after learning. Vertical bars show 90% confidence intervals. **b**, The y components of hand stiffness in the NF (green) and in the DF (red).

There was a selective increase in stiffness in the direction of instability. After learning, the stiffness of the arm increased sufficiently that the combined stiffness of the arm and the environment in the x direction (Fig. 4a, Net, black data) was positive and similar in magnitude to that measured during NF movements (Fig. 4a, NF, green data). Trial-to-trial variations in x force did not correlate with changes in stiffness in the x direction for any subject (Table 1), providing even stronger support for impedance control. Furthermore, we can conclude that the adaptation was preprogrammed and not reactive to the unstable environment because the DF was not applied on those trials used for stiffness measurement. (This, however, was not detected by the subjects.)

Previous impedance adaptation studies did not address the capacity to modify stiffness geometry independently of force. Adaptive increases in stiffness observed for single-joint tasks in unstable environments^{17,18,22} provide a limited view of the adaptive capabilities of the central nervous system (CNS) because the only parameter that can be varied is stiffness magnitude. Changes in stiffness geometry, observed in ball catching²⁰, occurred principally while joint torque was changing to compensate for the momentum of the ball. The anticipatory increase in stiffness just before ball impact involved generalized muscle coactivation, which serves mainly to increase stiffness magnitude. The use of biarticular shoulder muscles to perform a postural task, requiring no net shoulder torque⁹, might be considered a form of impedance control. However, because joint stiffness clearly scaled with joint torque in this task it is not clear whether the involvement of biarticular muscles represents genuine impedance control or ‘hard-wired’ muscle synergy. Similarly, the statistically significant covariation of change in stiffness with change in external force found for movements in a velocity-dependent force field²³ could be attributed largely to the linear relationship between joint torque and stiffness^{14,15}. The modifications of stiffness observed during movements constrained by a mechanical channel²⁴ were also probably due to changes in force. In contrast, adaptation to the DF required no change in force or joint torque relative to NF movements, yet subjects skilfully adapted their stiffness to the instability of the environment so that stiffness increased only in the required direction and by the required amount. Using the DF, we have been able to show for the first time that the CNS can voluntarily control the magnitude, shape and orientation of the endpoint stiffness in a predictive way that is independent of the force needed to compensate for the imposed dynamics.

Table 1 Correlation between stiffness and force

Subject	S1	S2	S3	S4	S5
R^2	0.24	0.53	0.08	0.29	0.34

R^2 is the correlation between the change in stiffness and the change in force in the DF relative to the NF. DF, divergent force field; NF, null field.

Finally, why does the CNS use the difficult and computationally costly strategy of controlling the full impedance geometry rather than the simple strategy of co-contracting all muscles? The particular shape of endpoint stiffness learned in the DF, large in the direction of instability and small in the direction of movement, suggests a trade-off between stability and high muscle activation. Increasing impedance enhances the robustness to external perturbations, but co-contraction of muscles increases metabolic cost¹². Moreover, motor output variability increases with muscle activation^{25,26}. Therefore, reducing the impedance will decrease both the metabolic cost and movement error. For these reasons, it is desirable to optimize impedance. The control problem faced by the CNS is to optimize the magnitude, shape and orientation of impedance to achieve stability while minimizing metabolic cost. □

Methods

Nine subjects performed the experiment (24–34 years of age; two females and seven males). The learning analysis used only data from the five naive subjects. The institutional ethics committee approved the experiments and the subjects gave informed consent before participation. The robot used is a parallel-link direct-drive air and magnet floating manipulandum (PFM), which moves in the horizontal plane. Details of its design and operation have been described previously^{19,21}. The PFM was set up by T. Yoshioka and H. Gomi. Subjects sat in a chair with a harness to constrain the trunk so that the elbow and shoulder joints could only move in the horizontal plane (two degrees of freedom). The forearm and wrist were held in a thermoplastic splint rigidly attached to the manipulandum. Therefore, subjects could not vary the kinematics of the arm to change its impedance. Subjects began by performing 100 movements in the NF, followed by 100–300 learning trials in the DF. Subjects then made 100 more movements in the DF, 20 of which had the force field removed unexpectedly (called after-effects). All movements were recorded during these sessions, including those not reaching the target. The start position, target (diameter 2.5 cm) and instantaneous hand position were displayed on a horizontal screen slightly above the arm. The prescribed movement time of 600 (±100) ms was indicated by acoustic signals.

Adaptation to the DF was quantified by calculating the error relative to a straight line joining the start position and target centre. The hand-path error *E* represents the area between the actual movement path and the straight line, and was calculated from the start time 0 (75 ms before crossing a hand-velocity threshold of 0.05 m s⁻¹), to the termination time *T* (when curvature exceeded 0.07 m⁻¹):

$$E = \int_0^T |x(t)||y'(t)|dt \tag{2}$$

To test whether learning occurred, we compared the hand-path error in the first five learning trials and the last five learning trials using an analysis of variance (ANOVA) with random factor subjects. An ANOVA was also used to compare the last five trials in the NF and the last five trials of the learning phase. The presence of after-effects was tested by performing an ANOVA comparing the hand-path error of the last 20 movements in the NF to 20 after-effects trials.

Before each stiffness measurement session, the subjects were retrained in the force field to ensure that they had remembered the field (NF, 40 trials; DF, 80 trials). Eighty movements were then recorded in the force field, of which 40, selected randomly, incorporated displacements. These displacements, at the midpoint of each movement, were randomly applied in one of eight directions (0°, 45°, 90°, 135°, 180°, 225°, 270° or 315°). The endpoint stiffness was estimated from the resulting mean change in hand force and position over a 60-ms interval during the perturbation, starting from 120 ms after onset of the perturbation¹⁹.

Linear regression was carried out to determine whether the change in stiffness for movements in the DF relative to those in the NF could be explained by the corresponding change in force. By definition, the difference between torque produced in response to a perturbation [*dq_s* *dq_e*]^T in the DF and NF is given by:

$$\begin{bmatrix} \delta\tau_s \\ \delta\tau_e \end{bmatrix} = \Delta\mathbf{K}_q \begin{bmatrix} dq_s \\ dq_e \end{bmatrix} \tag{3}$$

where $\Delta\mathbf{K}_q$ is the difference between DF and NF joint stiffness, and $\delta\tau_s$ and $\delta\tau_e$ correspond to the difference between DF and NF shoulder and elbow torque owing to elastic resistance λ , respectively.

It has been shown that the joint stiffness is linearly dependent on the joint torque in posture¹⁵. Making the same assumption for movement, we obtain the linear relation:

$$\Delta\mathbf{K}_q = \begin{bmatrix} \alpha_1\Delta\tau_s & \alpha_2\Delta\tau_e \\ \alpha_3\Delta\tau_e & \alpha_4\Delta\tau_e \end{bmatrix} \tag{4}$$

where $\alpha_1, \dots, \alpha_4$ are constants, and $\Delta\tau_s$ and $\Delta\tau_e$ are the difference between DF and NF shoulder and elbow driving torque measured by the force sensor λ . The difference in endpoint force can be obtained from equation (3) using the jacobian. Restricting the analysis to the *x* direction, since endpoint stiffness in the DF was principally modified in that direction, we then have:

$$\Delta F_x = [\Delta\tau_s dq_s \quad \Delta\tau_e dq_e \quad \Delta\tau_e dq_s] \begin{bmatrix} \gamma_1 \\ \gamma_2 \\ \gamma_3 \end{bmatrix} \tag{5}$$

where γ_1, γ_2 and γ_3 are constants that incorporate the α_i and the coefficients of the jacobian. R^2 values were computed for each subject using a linear system of 40 equations (with γ_1, γ_2 and γ_3 as unknowns) representing the 40 trials where stiffness was measured. An R^2 value close to one would indicate that endpoint stiffness in the *x* direction was completely determined by the joint torques required to perform the task. An R^2 value close to zero would indicate that in every trial endpoint stiffness was controlled independently of the required joint torques. A more detailed derivation of these equations is available as Supplementary Information.

Received 8 June; accepted 27 September 2001.

1. Shadmehr, R. & Mussa-Ivaldi, F. A. Adaptive representative of dynamics during learning of a motor task. *J. Neurosci.* **14**, 3208–3224 (1994).
2. Lackner, J. R. & Dizio, P. Rapid adaptation to Coriolis force perturbations of arm trajectory. *J. Neurophysiol.* **72**, 299–313 (1994).
3. Shadmehr, R. & Holcomb, H. H. Neural correlates of motor memory consolidation. *Science* **277**, 821–825 (1997).
4. Krakauer, J. W. Ghilardi, M. F. & Ghez, C. Independent learning of internal models for kinematic and dynamic control of reaching. *Nature Neurosci.* **2**, 1026–1031 (1999).
5. Kawato, M. Internal models for motor control and trajectory planning. *Curr. Opin. Neurobiol.* **9**, 718–727 (1999).
6. Thoroughman, K. A. & Shadmehr, R. Learning of action through adaptive combination of motor primitives. *Nature* **407**, 742–747 (2000).
7. Wolpert, D. M., Miall, R. C. & Kawato, M. Internal models in the cerebellum. *Trends Cogn. Sci.* **2**, 338–347 (1998).
8. Colgate, J. E. & Hogan, N. Robust control of dynamically interacting systems. *Int. J. Control.* **48**, 65–88 (1988).
9. McIntyre, J., Mussa-Ivaldi, F. A. & Bizzi, E. The control of stable postures in the multijoint arm. *Exp. Brain Res.* **110**, 248–264 (1996).
10. Rack, P. M. H. in *Handbook of Physiology* Sect. 1. *The Nervous System* (ed. Brooks, V. B.) 229–256 (American Physiological Society, Bethesda, 1981).
11. Milner, T. E. & Cloutier, C. Compensation for mechanically unstable loading in voluntary wrist movement. *Exp. Brain Res.* **94**, 522–532 (1993).
12. Hogan, N. The mechanics of multi-joint posture and movement control. *Biol. Cybern.* **52**, 315–331 (1985).
13. Won, J. & Hogan, N. Stability properties of human reaching movements. *Exp. Brain Res.* **107**, 125–136 (1995).
14. Hunter, I. W. & Kearney, R. E. Dynamics of human ankle stiffness: variation with mean ankle torque. *J. Biomech.* **15**, 747–752 (1982).
15. Gomi, H. & Osu, R. Task-dependent viscoelasticity of human multijoint arm and its spatial characteristics for interaction with environments. *J. Neurosci.* **18**, 8965–8978 (1998).
16. Mussa-Ivaldi, F. A., Hogan, N. & Bizzi, E. Neural, mechanical, and geometric factors subserving arm posture in humans. *J. Neurosci.* **5**, 2732–2743 (1985).
17. Akazawa, K., Milner, T. E. & Stein, R. B. Modulation of reflex EMG and stiffness in response to stretch of human finger muscle. *J. Neurophysiol.* **49**, 16–27 (1983).
18. De Serres, S. J. & Milner, T. E. Wrist muscle activation patterns and stiffness associated with stable and unstable mechanical loads. *Exp. Brain Res.* **86**, 451–458 (1991).
19. Burdet, E. et al. A method for measuring endpoint stiffness during multi-joint arm movements. *J. Biomech.* **33**, 1705–1709 (2000).
20. Lacquaniti, F., Carrozzo, M. & Borghese, N. A. Time-varying mechanical behavior of multijointed arm in man. *J. Neurophysiol.* **69**, 1443–1464 (1993).
21. Gomi, H. & Kawato, M. Human arm stiffness and equilibrium-point trajectory during multi-joint movement. *Biol. Cybern.* **76**, 163–171 (1997).
22. Milner, T. E., Cloutier, C., Leger, A. B. & Franklin, D. W. Inability to activate muscles maximally during co-contraction and the effect on joint stiffness. *Exp. Brain Res.* **107**, 293–305 (1995).
23. Burdet, E., Osu, R., Franklin, D. W., Milner, T. E. & Kawato, M. in *Proc. 1999 ASME Annu. Symp. on Haptic Interfaces and Virtual Environments for Teleoperator Systems (DSC7C-7)* 421–428 (ASME Press, New York, 1999).
24. Gomi, H. & Kawato, M. Task-dependent stiffness of human multi-joint arm during point-to-point movement. Technical Report ISRL-95-4 (NTT Basic Research Laboratories, Atsugi, 1995).
25. Schmidt, R. A., Zelaznik, H., Hawkins, B., Frank, J. S. & Quinn, J. T. Jr Motor-output variability: a theory for the accuracy of rapid motor acts. *Psychol. Rev.* **47**, 415–451 (1979).
26. Sliifkin, A. B. & Newell, K. M. Noise, information transmission, and force variability. *J. Exp. Psychol. Hum. Percept. Perform.* **25**, 837–851 (1999).

Supplementary Information accompanies the paper on Nature's website (<http://www.nature.com>).

Acknowledgements

The experiments were performed at ATR. We thank A. Illiesch and T. Flash for discussions. We thank A. Smith, S. Scott, D. Wolpert, Z. Ghahramani, and S. Schaal for giving valuable comments on an earlier version of the manuscript. This research was supported by the Special Coordination Fund for promoting Science and Technology of the Science and Technology agency of the Japanese Government; the Swiss National Science Foundation; the Natural Sciences and Engineering Research Council of Canada; and the Human Frontier Science Program.

Correspondence and requests for materials should be addressed to M.K. (e-mail: kawato@atr.co.jp). Further information about this work can be found at <http://www.bioeng.nus.edu.sg/humanRobotics.htm>.