

14. The differential response of excitatory and inhibitory cells can be justified physiologically. (i) As a function of the injected current, the firing rate in inhibitory cells increases twice as fast as that in pyramidal cells [D. A. McCormick, B. W. Connors, J. W. Lighthall, D. A. Prince, *J. Neurophysiol.* **54**, 782 (1985)]. (ii) Inhibitory cells in vitro adapt only weakly (*ibid.*). (iii) Successive excitatory post-synaptic potentials (EPSPs) from pyramidal cells onto inhibitory interneurons are potentiated, whereas successive EPSPs from pyramidal cells onto pyramidal cells are depressed [A. M. Thomson and D. C. West, *Neurosci.* **54**, 329 (1993); A. M. Thomson, J. Deuchars, D. C. West, *ibid.*, p. 347]. (iv) Long-range connections terminate in the more distal regions of the dendritic tree of pyramidal cells [L. J. Cauler and B. W. Connors, in *Single Neuron Computation*, T. McKenna, J. Davis, Z. Zornetzer, Eds. (Academic Press, San Diego, 1992), pp. 199–299], and conductances in these regions of high dendritic input resistance are more likely to saturate, or to be shunted [Ö. Bernander, C. Koch, R. J. Douglas, *J. Neurophysiol.* **72**, 2743 (1994)].
15. J. A. Hirsch and C. D. Gilbert, *J. Neurosci.* **11**, 1800 (1991).
16. K. Wiesenfeld and F. Moss, *Nature* **373**, 33 (1995).
17. Signal detection depends on the discrimination between the signal plus noise and noise alone. To first order, the probability of correct classification is proportional to the difference in firing rates $f(I) - f(0)$, where $f(I)$ is the firing rate of the cell in the presence of a center stimulus, and $f(0)$ is the firing rate in its absence. To second order, detection depends on the probability distribution of spike counts over a finite time window. Given a near-threshold center stimulus and a fixed observation time window, the detection probability will peak at an optimal noise level; this peak is termed the “stochastic resonance.”
18. For an integrate-and-fire cell, additional input variance results in an increased response to weak subthreshold stimuli, but only a negligible effect on strong stimuli. Integrate-and-fire models capture many of the key features of the discharge curves of real cells in response to noisy current injections [Z. F. Mainen and T. J. Sejnowski, *Science* **268**, 1503 (1995); A. Zador, personal communication]. In these studies, it has been shown that increasing the input variance linearizes the discharge response curve at low firing rates, so that the cell fires even in response to subthreshold input. The fluctuations in the lateral cortico-cortical input current are significant because of the irregularity in the spiking of surround cells. If the surround provides a signal of constant variance and small negative mean, at low stimulus contrast the variance effect will dominate and lower the threshold for detection, whereas at high stimulus contrast the negative mean current will result in the suppression of redundant information in a high-contrast texture background.
19. M. Cannon and S. Fullenkamp, *Vision Res.* **33**, 1685 (1993).
20. D. C. Somers *et al.* have independently developed a model similar to ours [D. Somers, S. Nelson, M. Sur, *Soc. Neurosci. Abstr.* **20**, 1577 (1994)] and verified some of these predictions experimentally [D. C. Somers *et al.*, in *Lateral Interactions in the Cortex*, J. Sirosh and R. Miikkulainen, Eds. (University of Texas, Austin, in press)].
21. Supported by the Howard Hughes Medical Institute, the National Institutes of Mental Health (grants MH47566 and MH45156), the Office of Naval Research, the Air Force Office of Scientific Research, the National Science Foundation, the Center for Neuromorphic Systems Engineering as a part of the National Science Foundation Engineering Research Center Program, and by the Office of Strategic Technology of the California Trade and Commerce Agency. We thank J. Knerim, K. Grieve, F. Wörgötter, and C. Koch for discussions; A. Zador, U. Polat, and M. Sur for access to unpublished data; G. Blasdel for providing the orientation map underlying Fig. 1; and C. Koch and J. McClelland for a stimulating work environment.

3 April 1995; accepted 27 July 1995

An Internal Model for Sensorimotor Integration

Daniel M. Wolpert,* Zoubin Ghahramani, Michael I. Jordan

On the basis of computational studies it has been proposed that the central nervous system internally simulates the dynamic behavior of the motor system in planning, control, and learning; the existence and use of such an internal model is still under debate. A sensorimotor integration task was investigated in which participants estimated the location of one of their hands at the end of movements made in the dark and under externally imposed forces. The temporal propagation of errors in this task was analyzed within the theoretical framework of optimal state estimation. These results provide direct support for the existence of an internal model.

The notion of an internal model, a system that mimics the behavior of a natural process, has emerged as an important theoretical concept in motor control (1). There are two varieties of the internal model: (i) forward models, which mimic the causal flow of a process by predicting its next state (for example, position and velocity) given the current state and the motor command; and (ii) inverse models, which invert the causal flow by estimating the motor command that caused a particular state transition. Forward models have been shown to be of potential use for solving four fundamental problems in computational motor control. First, the delays in most sensorimotor loops are large, making feedback control too slow for rapid movements. With the use of a forward model for internal feedback, the outcome of an action can be estimated and used before sensory feedback is available (2, 3). Second, a forward model is a key ingredient in a system that uses motor outflow (also called efference copy) to anticipate and cancel the sensory effects of movement (also called refference) (4). Third, a forward model can be used to transform errors between the desired and actual sensory outcome of a movement into the corresponding errors in the motor command, thereby providing appropriate signals for motor learning (5). Similarly, by predicting the sensory outcome of the action without actually performing it, a forward model can be used in mental practice to learn to select between possible actions (6). Finally, a forward model can be used for state estimation in which the model's prediction of the next state is combined with a reafferent sensory correction (7). Although shown to be of theoretical use, the existence of an internal forward model in the central nervous system (CNS) is still a topic of debate.

When we move an arm in the absence of visual feedback, there are three basic methods the CNS can use to obtain an estimate of the current state—the position and velocity—of the hand. The system can make use of sensory inflow (the information available from proprioception), it can make use of integrated motor outflow (the motor commands sent to the arm), or it can combine these two sources of information by use of a forward model. To test between these possibilities, we carried out an experiment in which participants, after initially viewing one of their arms in the light, made arm movements in the dark. Three experimental conditions were studied, involving the use of null, assistive, and resistive force fields. We assessed the participants' internal estimate of hand location by asking them to localize visually the position of their hand at the end of the movement (8). The bias of this location estimate, plotted as a function of movement duration, shows a consistent overestimation of the distance moved (Fig. 1). This bias shows two distinct phases as a function of movement duration: an initial increase reaching a peak of 0.9 cm after 1 s followed by a sharp transition to a region of gradual decline. The variance of the estimate also shows an initial increase during the first second of movement after which it plateaus at about 2 cm². External forces had distinct effects on the bias and variance propagation. Whereas the bias was increased by the assistive force and decreased by the resistive force, the variance was unaffected.

These experimental results can be fully accounted for if we assume that the motor control system integrates the efferent outflow and the reafferent sensory inflow. To establish this conclusion, we developed an explicit model of the sensorimotor integration process, which contains as special cases all three of the methods referred to above (9). This model is based on the observer framework (7) from engineering in which the state estimator (or observer) has access to both the inputs and outputs of the system. Specifically, the input to the arm is the

Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.

*Present address to which correspondence should be addressed: Sobell Department of Neurophysiology, Institute of Neurology, Queen Square, London WC1N 3BG, UK.

motor command and the output is the sensory feedback that, in the absence of vision, consists solely of proprioception. On the basis of these two sources, the observer produces an estimate of the state of the system. In particular, we chose to use a Kalman filter (10) observer, which is a linear dynamical system that produces an estimate of the location of the hand by using both the motor outflow and sensory feedback in conjunction with a model of the motor system. Using these sources of information, the model estimates the arm's state, integrating sensory

and motor signals to reduce the overall uncertainty in its estimate.

The model is a combination of two processes that together contribute to the state estimate. The first process (upper part, Fig. 2A) uses the current state estimate and motor command to predict the next state by simulating the movement dynamics with a forward model. The second process (lower part, Fig. 2A) uses a model of the sensory output process to predict the sensory feedback from the current state estimate. The sensory error—the difference between actual

and predicted sensory feedback—is used to correct the state estimate resulting from the forward model. The relative contributions of the internal simulation and sensory correction processes to the final estimate are modulated by the Kalman gain so as to provide optimal state estimates. By making particular choices for the parameters of the Kalman filter, we were able to simulate motor outflow-based estimation (11), sensory inflow-based estimation, and forward model-based sensorimotor integration. Moreover, to accommodate the observation that participants generally tend to overestimate the distance that their arm has moved, we set the gain that couples force to state estimates to a value that is larger than its veridical value (12). All other components of the internal model were set to their veridical values.

The Kalman filter model demonstrates the two distinct phases of bias propagation observed (Fig. 2, B through E). By overestimating the force acting on the arm, the forward model overestimates the distance traveled, an integrative process eventually balanced by the sensory correction. The model also captures the differential effects on bias of the externally imposed forces. By overestimating an increased force under the assistive condition, the bias in the forward model accrues more rapidly and is balanced by the sensory feedback at a higher level. The converse applies to the resistive force. The pattern of variance propagation is also captured by the model. The variance of the state estimate derives from two sources of variance in the system: the first is the variability in the response of the arm to the motor commands and the second is the noise in the subsequent sensory feedback. Initially, when the hand is in view, the state estimate is assumed to be accurate. The accuracy of the prediction from the forward model component of the Kalman filter depends on the accuracy of the current state estimate (one of its inputs). Therefore, during the early part of the movement, when the current state estimate is accurate, the sensorimotor integration process weights heavily the contribution of the forward model to the final estimate. However, in the later stages of the movement, when the current state estimate is less accurate, the sensory feedback must be relied on to correct for inaccuracies in the forward model. In the Kalman filter, the relative weighting shifts from the forward model toward sensory feedback over the first second of movement and then remains approximately constant in the asymptote of the variance propagation. In accord with the experimental results, the model predicts no change in variance under the two force conditions.

We have shown that the Kalman filter is able to reproduce the propagation of the

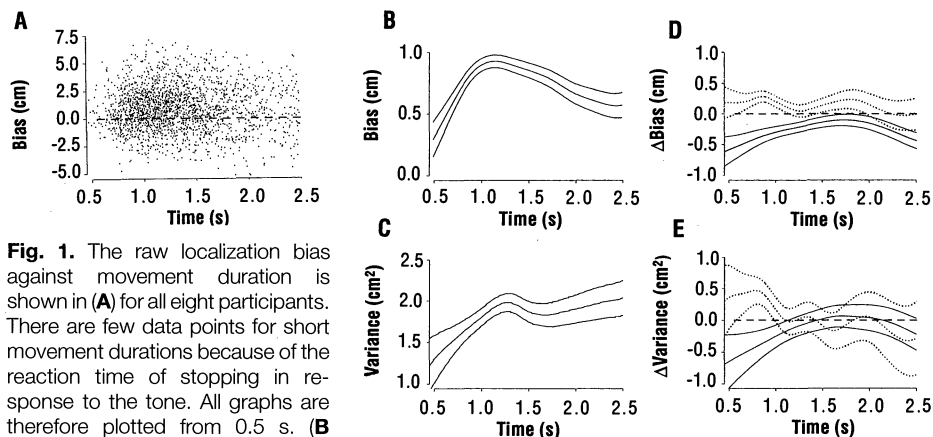
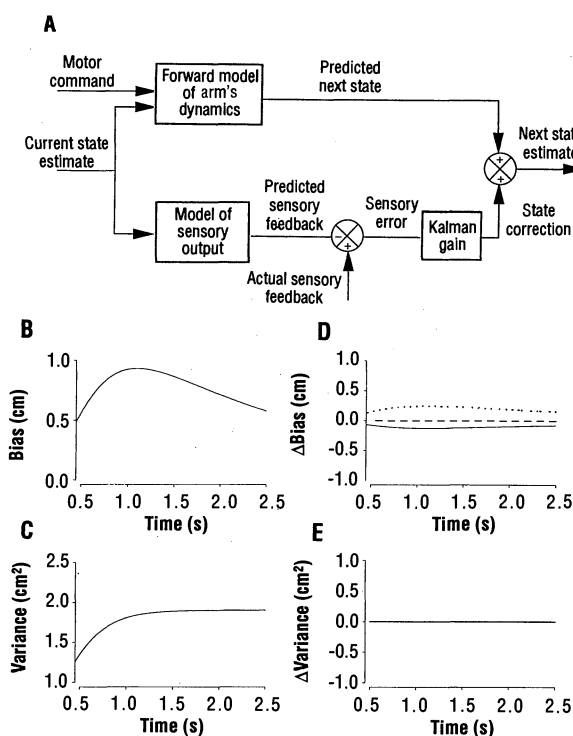


Fig. 1. The raw localization bias against movement duration is shown in (A) for all eight participants. There are few data points for short movement durations because of the reaction time of stopping in response to the tone. All graphs are therefore plotted from 0.5 s. (B through E) The main effect fits of the generalized additive model to the data. The propagation of the (B) bias and (C) variance of the state estimate is shown, with outer standard error lines, against movement duration. The differential effects on (D) bias and (E) variance of the external force, assistive (dotted lines) and resistive (solid lines), are also shown relative to zero (dashed line). A positive bias represents an overestimation of the distance moved. The difference in variance propagation between the resistive and assistive fields was not significant over the movement; the difference in bias was significant at the $P = 0.05$ level.

Fig. 2. (A) The Kalman filter model is shown schematically, consisting of two processes. The first (upper part) uses the motor command and the current state estimate to achieve a state estimate using the forward model to simulate the arm's dynamics. The second process (lower part) uses the difference between expected and actual sensory feedback to correct the forward model state estimate. The relative weighting of these two processes is mediated through the Kalman gain. (B through E) Simulated bias and variance propagation, in the same representation and scale as Fig. 1, B through E, from the Kalman filter model of the sensorimotor integration process.



bias and variance of estimated position of the hand as a function of both movement duration and external forces. The Kalman filter model suggests that the peaking and gradual decline in bias is a consequence of a trade-off between the inaccuracies accumulating in the internal simulation of the arm's dynamics and the feedback of actual sensory information. Simple models that do not trade off the contributions of a forward model with sensory feedback, such as those based purely on sensory inflow or on motor outflow, are unable to reproduce the observed pattern of bias and variance propagation (13). The ability of the Kalman filter to parsimoniously model our data suggests that the processes embodied in the filter—namely, internal simulation through a forward model together with sensory correction—are likely to be embodied in the sensorimotor integration process. We feel that the results of this state estimation study provide evidence that a forward model is used by the CNS in maintaining its estimate of the hand location. Furthermore, the state estimation paradigm provides a framework to study the sensorimotor integration process in both normal and patient populations. The model predicts monotonically increasing bias and variance, if the afferent signal is eliminated, and under-shoot rather than overshoot in bias propagation if the forward model is eliminated. These specific predictions can be tested in both patients with sensory neuropathies, who lack proprioceptive reafference, and patients with damage to the cerebellum, a proposed site for the forward model (3).

REFERENCES AND NOTES

- M. I. Jordan, in *Handbook of Perception and Action: Motor Skills*, H. Heuer and S. Keele, Eds. (Academic Press, New York, 1995); M. Kawato, K. Furawaka, R. Suzuki, *Biol. Cybern.* **56**, 1 (1987).
- M. Ito, *The Cerebellum and Neural Control* (Raven, New York, 1984).
- R. C. Miall, D. Weir, D. M. Wolpert, J. Stein, *J. Mot. Behav.* **25** (3), 203 (1993).
- C. Gallistel, *The Organization of Action: A New Synthesis* (Erlbaum, Hillsdale, NJ, 1980); D. Robinson, J. Gordon, S. Gordon, *Biol. Cybern.* **55**, 43 (1986).
- M. I. Jordan and D. Rumelhart, *Cognit. Sci.* **16**, 307 (1992).
- R. Sutton and A. Barto, *Psychol. Rev.* **88**, 135 (1981).
- G. Goodwin and K. Sin, *Adaptive Filtering Prediction and Control* (Prentice-Hall, Englewood Cliffs, NJ, 1984).
- The experimental setup consisted of a planar, virtual visual feedback system [described in D. M. Wolpert, Z. Ghahramani, M. I. Jordan, *Exp. Brain Res.* **103**, 460 (1995)] in conjunction with a planar, two degree-of-freedom manipulandum driven by two torque motors [described in I. Faye, thesis, Massachusetts Institute of Technology (1986)]. Each participant gripped a manipulandum on which his thumb was mounted. The manipulandum was used to accurately measure the position of the participant's thumb and also, using the torque motors, to apply forces to the hand. The hand was constrained to move along a straight line passing transversely in front of the participant. The virtual visual feedback system was used to project computer-controlled images into the plane of the movement. Eight untrained male participants, who gave their informed consent, performed 300 trials each. Each trial started with the participant visually placing his thumb at a target square projected randomly on the movement line. The arm was then illuminated for 2 s, thereby allowing the participant to perceive visually his initial arm configuration. The light was then extinguished, leaving just the initial target. The participant was then required to move his hand slowly either to the left or right, as indicated by an arrow in the initial starting square. This movement was made in the absence of any visual feedback of the participant's arm configuration. The participant was instructed to move until he heard a tone, at which point he stopped. The timing of the tone was controlled to produce a uniform distribution of path lengths from 0 to 30 cm. During this movement, the participant moved either in a randomly selected null or constant assistive or resistive force field of 3 N generated by the torque motors. Although it is not possible to directly probe a participant's internal representation of the state of his arm, we examined a function of this state: the estimated visual location of the thumb. The relation between the state of the arm and the visual coordinates of the hand is known as the kinematic transformation [J. Craig, *Introduction to Robotics* (Addison-Wesley, Reading, MA, 1986)]. Therefore, once at rest the participant indicated the visual estimate of the unseen thumb position using a trackball, held in his other hand, to move a cursor projected in the plane of the thumb along the movement line. The discrepancy between the actual and visual estimate of thumb location was recorded as a measure of the state estimation error. The bias and variance propagation of the state estimate were analyzed as a function of movement duration and external forces. A generalized additive model (GAM) [T. Hastie and R. Tibshirani, *Generalized Additive Models* (Chapman and Hall, London, 1990)] with smoothing splines (five effective degrees of freedom) was fit to the bias and variance as a function of final position, movement duration, and the interaction of the two forces with movement duration, simultaneously for main effects and for each participant. Errors related to the final position factor represent movement-independent inaccuracies in visually locating the hand and can be attributed to the kinematic transformation, which relates the state estimate of hand position to the perceived visual location. As these static distortions of the kinematic transformation are not relevant to our study of movement-related errors, they were factored out by the GAM fit. Although distance moved was correlated with movement duration ($r^2 = 0.35$), its inclusion as an additional factor in the model had a minimal effect on the component fits of duration and external force.
- The system dynamics of the hand were approximated by a damped (coefficient β) point mass m , moving in one dimension acted on by a force $u = u_{\text{int}} + u_{\text{ext}}$, combining both internal (int) motor commands and external (ext) forces. Representing the state of the hand at time t as $\mathbf{x}(t)$ (a 2×1 vector of position and velocity) and its time derivative by $\dot{\mathbf{x}}(t)$, the system dynamic equations can be written in the general form of

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}u(t) + \mathbf{w}(t)$$
 where

$$\mathbf{A} = \begin{bmatrix} 0 & 1 \\ 0 & -\beta/m \end{bmatrix} \text{ and } \mathbf{B} = \begin{bmatrix} 0 \\ 1/m \end{bmatrix}$$
 and the vector $\mathbf{w}(t)$ represents the process of white noise with an associated covariance matrix given by $\mathbf{Q} = E[\mathbf{w}(t)\mathbf{w}(t)^T]$. The system has an observable output, the sensory information, representing the proprioceptive signals (for example, from muscle spindles and joint receptors). This output, $\mathbf{y}(t)$, is linked to the actual hidden state $\mathbf{x}(t)$ by $\mathbf{y}(t) = \mathbf{C}\mathbf{x}(t) + \mathbf{v}(t)$, where the vector $\mathbf{v}(t)$ represents the output white noise, which has the associated covariance matrix $\mathbf{R} = E[\mathbf{v}(t)\mathbf{v}(t)^T]$. We represent the internal estimate of the state at time t by $\hat{\mathbf{x}}(t)$. We assumed that this system is fully observable and chose \mathbf{C} to be the identity matrix. At time $t = 0$, the participant was given full view of his arm and, therefore, started with an estimate $\hat{\mathbf{x}}(0) = \mathbf{x}(0)$ with zero bias and variance; we assumed that vision calibrates the system. At this time, the light was extinguished and the participant had to rely on the inputs and outputs to estimate the system's state. Using a model of the system $\hat{\mathbf{A}}$, $\hat{\mathbf{B}}$, and $\hat{\mathbf{C}}$, the Kalman filter provides an optimal linear estimator of the state given by

$$\hat{\mathbf{x}}(t) = \underbrace{\hat{\mathbf{A}}\hat{\mathbf{x}}(t) + \hat{\mathbf{B}}u(t)}_{\text{Forward model}} + \underbrace{K(t)[\mathbf{y}(t) - \hat{\mathbf{C}}\hat{\mathbf{x}}(t)]}_{\text{Sensory correction}}$$
 where $K(t)$ is the recursively updated gain matrix. This state estimate combines an estimate from the internal model of the system dynamics together with a sensory correction modulated by the Kalman gain matrix $K(t)$. We used this state update equation to model the bias and variance propagation and the effects of the external force. The parameters in the simulation, $\beta = 3.9 \text{ N}\cdot\text{s m}^{-1}$, $m = 4 \text{ kg}$, and u were chosen on the basis of the mass of the arm and the observed relation between time and distance traveled. Specifically, the total force u was chosen to be linearly related to the average velocity under each of the three force conditions: 1.3, 1.5, and 1.9 N, corresponding to the average movement velocities of 10.8, 12.8, and 16.6 cm s^{-1} for the resistive, null, and assistive conditions, respectively. To end the movement, the sign of the force acting on the hand was reversed until the arm was stationary. To simulate the overestimation of distance traveled, $\hat{\mathbf{B}}$ was set to

$$\begin{bmatrix} 0 \\ 1.4/m \end{bmatrix}$$
 while both $\hat{\mathbf{A}}$ and $\hat{\mathbf{C}}$ accurately reflected the true system. Noise covariance matrices of $\mathbf{Q} = 9.5 \times 10^{-5}$ and $\mathbf{R} = 3.3 \times 10^{-4}$ were used, where I is the identity matrix. This represents a standard deviation of 1.0 cm for the position output noise and 1.8 cm s^{-1} for the position component of the state noise. The parameters $\hat{\mathbf{B}}$, \mathbf{Q} , and \mathbf{R} were chosen by trial and error to show that this model is able to qualitatively capture the data. The increasing and plateauing nature of the variance was robust to changes in these parameters. This behavior is a basic feature of Kalman filter models and was observed in all the simulations run. The rate of rise and plateau level of the variance is determined by the relative variances \mathbf{Q} and \mathbf{R} . Provided $\hat{\mathbf{B}}$ is chosen to be larger than the true value of \mathbf{B} , the bias shows a typical increasing phase followed by a slow decline. The exact rate of rise and peak of the bias depend on the particular choice of $\hat{\mathbf{B}}$. However, in general, as long as the basic components of the Kalman filter are maintained the shapes of the simulation plots are not particularly sensitive to the choice of parameters.
- R. Kalman and R. S. Bucy, *J. Basic Eng.* **83D**, 95 (1961).
- Estimation based purely on motor outflow is also known as "dead reckoning." This uses the rate of change of a variable, as estimated by a forward model, to update the current estimate. This term derives from its usage by sailors in navigation, who would estimate the position of their ship at sea on the basis of their previous position, time elapsed, and their estimated velocity over the ground. By effectively internally modeling the ship's dynamics, the sailors would learn to estimate the velocity based on the observed heading of the ship, the sails set, the force and wind direction, and the currents. For a review of dead reckoning in animal behavior, see C. Gallistel, *The Organization of Learning* (MIT Press, Cambridge, MA, 1990), chap. 4.
- This setting is consistent with the independent data that participants tend to under-reach in pointing tasks, which suggests an overestimation of distance traveled [J. Soechting and M. Flanders *J. Neurophysiol.* **62**, 582 (1989)].
- A model based purely on motor outflow (dead reckoning) produces a monotonically increasing bias and variance. Models based purely on sensory inflow (reafference) cannot model the differential effects of the external forces on the bias propagation.
- We thank P. Dayan for suggestions about the manuscript. This project was supported by grants from the McDonnell-Pew Foundation, ATR Human Information Processing Research Laboratories, Siemens Corporation, and the U.S. Office of Naval Research. D.M.W. and Z.G. are McDonnell-Pew Fellows in Cognitive Neuroscience. M.I.J. is an NSF Presidential Young Investigator.

27 March 1995; accepted 7 August 1995