J Neurophysiol
87: 2434–2440, 2002; 10.1152/jn.00875.2001.

Controlling the Statistics of Action: Obstacle Avoidance

ANTONIA F. DE C. HAMILTON AND DANIEL M. WOLPERT
Sobell Department of Motor Neuroscience and Movement Disorders, Institute of Neurology, University College London, London WC1N 3BG, United Kingdom

Received 24 October 2001; accepted in final form 10 January 2002

Hamilton, Antonia F. de C. and Daniel M. Wolpert. Controlling the statistics of action: obstacle avoidance. *J Neurophysiol* 87: 2434–2440, 2002; 10.1152/jn.00875.2001. Task optimization in the presence of signal-dependent noise (TOPS) has been proposed as a general framework for planning goal-directed movements. Within this framework, the motor command is assumed to be corrupted by signal-dependent noise, which leads to a distribution of possible movements. A task can then be equated with optimizing some function of the statistics of this distribution. We found the optimal trajectory for obstacle avoidance by minimizing the mean-squared error at the end of the movement while keeping the probability of collision with the obstacle below a fixed limit. The optimal paths accurately predicted the empirical trajectories. This demonstrates that controlling the statistics of movements in the presence of signal-dependent noise may be a fundamental and unifying principle of goal-directed movements.

INTRODUCTION

Almost any motor task can in theory be achieved using an infinite variety of hand paths, velocity profiles, joint configurations, and muscle co-contraction levels. Yet, despite this redundancy in the motor system, almost every study has shown stereotypical patterns in human movement. Recently, a theoretical framework, task optimization in presence of signal-dependent noise (TOPS), has been proposed to account for the stereotypy in goal-directed movements (Harris and Wolpert 1998). Here we compare the performance of TOPS with other models of movement planning (Flash and Hogan 1985; Sabes and Jordan 1997; Uno et al. 1989) in an obstacle-avoidance task (Sabes et al. 1997, 1998).

The TOPS framework is based on optimal control, in which each possible movement is assigned a cost and the movement with the lowest cost is executed. Specifically, the TOPS framework proposes that motor commands are corrupted by signal-dependent noise, that is, noise whose SD increases linearly with the absolute level of the motor command (constant coefficient of variation). Such signal-dependent noise can be seen behaviorally in isometric tasks: when subjects are asked to generate either force pulses or constant levels of force, the SD of the force increases linearly with its mean level (Schmidt et al. 1979; Slifkin and Newell 1999). Similarly, the variability of the amplitude of a movement increases as the desired amplitude increases (Fitts 1966).

Given signal-dependent noise, the same desired motor command repeated many times leads to a distribution of possible states of the motor system, for example, a distribution of positions of the hand. TOPS proposes that selecting a movement for a given task can be equated with optimizing some function over this distribution, which leads to a unique optimal trajectory (although in theory it may be possible to have more than one optimal solution for some tasks). For example, by minimizing mean-squared endpoint error in a goal-directed eye or arm movement, TOPS is able to model the observed trajectories of both saccadic eye and point-to-point arm movements (Harris and Wolpert 1998). When subjects are asked to draw ellipses, the relationship between the curvature of the hand path and the hand velocity are related by a power law, known as the two-thirds power law (Viviani and Terzuolo 1982). This power law is predicted by the TOPS framework if the mean-squared deviation of the hand from the desired elliptical path is minimized (Harris and Wolpert 1998). Here we extend the TOPS model to a more complex task, that of obstacle avoidance, in which new stereotypical patterns have recently been observed (Sabes et al. 1997, 1998).

Subjects generally produce asymmetric paths when asked to move around an obstacle placed symmetrically between the start and the target (Fig. 1*A*). These paths can be characterized in terms of the near-point, the point on the path nearest to the obstacle (Fig. 1*B*). Sabes and Jordan (1997) examined how the near-point changes as the entire start, target, and obstacle were rotated about the obstacle tip. They sampled the obstacle orientations uniformly through 360°, but found that near-points from the set of trajectories were not uniformly distributed. In studies in both two (Sabes and Jordan 1997) and three dimensions (Sabes et al. 1998), Sabes et al. found that the near-points deviate from a symmetric distribution to cluster toward a preferred axis (see, for example, Fig. 3*A*) and that this preferred axis changed over the workspace.

Sabes and Jordan interpreted this anisotropic distribution in terms of inertial stability of the arm, as characterized by its mobility ellipse (Hogan 1985). They proposed that subjects skewed their paths to pass nearer the obstacle at a place where they were more stable to any perturbations that might cause a collision. They found good agreement between the empirically measured axis of the mobility ellipse and the preferred axis of the near-points. However, this explanation cannot be used either to predict the trajectories or to determine the extent to which the paths should be skewed.

Other current models of trajectory planning which could be

Address for reprint requests: A. Hamilton, Sobell Dept. of Motor Neuroscience and Movement Disorders, Institute of Neurology, Queen Square, London WC1N 3BG, UK (E-mail: a.hamilton@ion.ucl.ac.uk).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked "*advertisement*" in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

 0022-3077/02 $5.00 Copyright © 2002 The American Physiological Society

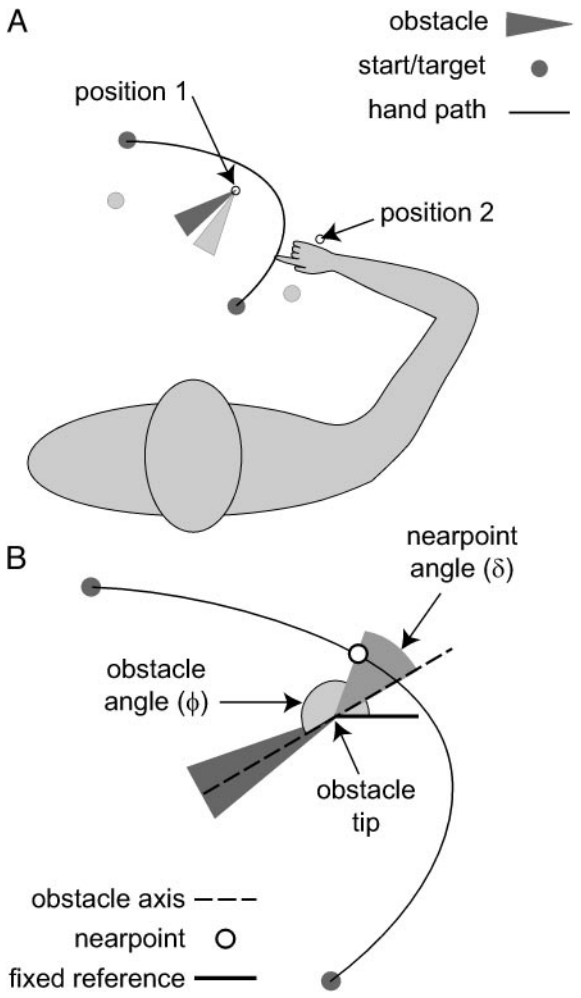

FIG. 1. *A*: schematic of obstacle avoidance task. Two positions in the workspace are shown. The dark gray circles and wedges indicate the start, target, and obstacle, which are enlarged in *B*. On successive trials, the obstacle tip remains in the same place, and the start, target, and obstacle rotate together about the tip so that all obstacle angles are sampled. The pale circles and wedges are an example of a different obstacle orientation. *B*: detail of the obstacle avoidance task. The dark gray wedge is the obstacle, and the dashed line is the obstacle axis. Start and target locations (dark gray circles) were 25 cm apart, and the obstacle was 8 cm long. The near-point (open circle) is the point on the path (thin solid line) closest to the obstacle tip. The obstacle angle (pale gray) is the angle between the obstacle axis and the 3 o'clock reference line (heavy black line). The near-point angle (mid gray) is the angle between the near-point and the obstacle axis.

applied to the obstacle avoidance task cannot predict the observed behavior. The minimum jerk model (Flash and Hogan 1985) proposes that movements of the hand are smooth and the cost is the integrated squared jerk (rate of change of acceleration) of the hand over the movement. These movements are planned only in endpoint space, independent of the arms dynamics, and therefore, the model predicts symmetric paths for every obstacle orientation and no preferred axis. The minimum torque change model (Uno et al. 1989) is also based on smoothness, but at the torque level. The cost is the integrated squared torque change summed over the joints and the movement. Both the minimum jerk model and the minimum torque change models predict that the trajectories would pass as close as possible to the obstacle for two reasons. First, the optimal paths without an obstacle are either straight (minimum jerk) or close to straight (minimum torque-change), and therefore, the cost reduces the nearer the path comes to the obstacle. Second, both models assume no variability in the motor command, so there is no penalty for coming extremely close to the obstacle.

Here we use the TOPS framework to simulate the obstacle avoidance task, by minimizing the mean-squared error at the target while ensuring that the probability of collision with the obstacle remains below a fixed limit. We assess the performance of the model against the performance of the mobility ellipse model proposed by Sabes and Jordan and show that TOPS is able to predict the trajectory, the amount of trajectory skewing, and the preferred axis.

METHODS

We simulated the obstacle avoidance task of Sabes and Jordan (1997), in which subjects made arm movements in the horizontal plane between two targets while avoiding a wedge-shaped obstacle. The start and target locations were 25 cm apart and the obstacle protruded 8 cm perpendicular from the direct line between the start and target points (Fig. 1*B*). For each trial, the obstacle orientation was selected from 180 different orientations, equally spaced at 2° intervals around the circle. The orientation determined both the obstacle angle and the positions of the targets so that the geometry of the task was preserved (apart from the rotation about the obstacle tip). The experiment was repeated at two positions in the workspace (Fig. 1*A*) for movements in both the clockwise and the counter-clockwise directions.

We simulated this task within the TOPS framework in which we assumed that motor commands are corrupted by signal-dependent noise and that the subjects choose the motor command which minimizes the mean-squared error at the target while limiting the probability of collision to below a fixed limit. Given this cost and this noise, the TOPS framework was used to predict the optimal feedforward trajectory, that is, with no on-line corrections for the noise, for a two-joint model of the arm.

To find the optimal path for a given obstacle orientation and location, a candidate path was constructed and the cost of this path evaluated. To parameterize the path, we used a quintic spline to specify the trajectory in Cartesian coordinates, x and y, each as a function of time. A quintic spline was chosen because it is smoothly differential up to fourth order, a necessary condition to generate a smoothly modulating motor command for the arm, which we modeled as a fourth-order system. Five knots, equally spaced in time, were used with the first and last knots placed on the start and target locations. The spline was constrained so that the velocity and acceleration were zero at the start and end of the movement (higher derivatives cannot be constrained with a quintic spline). The trajectory was therefore determined by six parameters, the x and y positions of the three interior knots. The movement duration was set at 736 ms (Sabes and Jordan 1997) and a sampling interval of 9.8 ms was chosen which defined 75 points along the path. A postmovement period of 500 ms (51 points) in which no motor command was generated was included to assess the error at the end of the movement.

To determine the motor command needed to achieve a candidate trajectory and the effects of signal-dependent noise on this command, the arm was modeled with two links moving in the horizontal plane (upper:lower link lengths, 0.315:0.459 m; center of mass, 0.13:0.15 m; mass, 0.9:1.1 kg; inertia, 0.0201:0.0453 kg/m^2; joint viscosity, 0.4 Nms/rad; joint stiffness was assumed to be dominated by stiffness of the muscles which is modeled separately) (Kawato 1996). Torques at each joint were generated by a linear second-order muscle, with time constants 30 and 40 ms (van der Helm and Rozendaal 2000). To determine the motor commands, the muscle model needs to possess *1*) the ability to invert and *2*) an order no higher than of two. With a higher order muscle and a quintic spline, the motor command would

become discontinuous because the overall order of the system would be greater than four (requiring more than 4 differentiations of the trajectory). We therefore chose a second-order linear muscle, as we have previously shown that using this simple representation we are able to model both the point-to-point arm trajectories and the two-thirds power law (Harris and Wolpert 1998).

Using inverse kinematics and dynamics, the motor command required to generate the candidate trajectory was determined. To evaluate the cost of the movement, the motor command (that is, the input to the muscle model) was corrupted by signal-dependent noise and the mean-squared error at the end of the movement was calculated. This was achieved using a new computationally efficient algorithm called the unscented transformation (Julier and Uhlmann 1995, 1996). The unscented transformation works by propagating the arm's state variables through the model of the arm for a particular noise distribution on the motor command and achieves an unbiased estimate of the mean and covariance of the path for every time step (see APPENDIX for details).

For each point along the trajectory, the covariance matrix $\mathbf{C}_t$, representing the variability in hand position at time t, was calculated. We then calculated a hit index $\mathbf{H}_t$ for each time step, which indicated whether the collision probability at that time step had been exceeded. This was achieved by constructing a covariance ellipses scaled to the appropriate collision probability (for example, for a collision probability of 2.3% a contour ellipse at 2 SDs was constructed) and testing if the ellipse intersected the obstacle ($\mathbf{H}_t = 1$) or not ($\mathbf{H}_t = 0$). The cost of a candidate path was taken as Trace $[\mathbf{C}_T] + K \sum_{t=1}^{T} \mathbf{H}_t$, where T is the movement duration. K was chosen as large enough (1000) so that after optimization $\mathbf{H}_t$ was zero for all time steps; that is, the collision probability was never exceeded. In other words, after optimization the cost was always Trace $[\mathbf{C}_T]$, that is, the mean-squared error of the final hand position.

The six free parameters, that is, the (x, y) coordinates of the three interior knots of the path, were optimized using simulated annealing (Kirkpatrick et al. 1983; Metropolis et al. 1953; Press et al. 1992). The initial setting of the interior knots placed the first and third free knots 16 cm apart and 2 cm beyond the obstacle tip and the second free knot 6 cm beyond the obstacle tip in line with the obstacle axis, so as to give a symmetric path that avoided the obstacle. Unlike gradient-descent methods which always aim to go "downhill," thereby reducing the cost at every cycle of the optimization, simulated annealing can allow the cost to increase, which helps the algorithm avoid local minima. Specifically, a change in the knots is accepted if the cost decreases or if the cost increase is less than the temperature parameter multiplied by a logarithmically distributed random number. As the simulated annealing progresses, the temperature gradually decreases, in our case by a factor of 10 every 100 cycles, where each cycle tests a set of seven candidate trajectories. The annealing was scheduled to restart using the optimal vertices and a Gaussian perturbation ($\mu = 0$, $\sigma = 0.05$ m) on each sub-optimal vertex for five optimizations in succession, because pilot simulations revealed that the final cost fell by a small amount when extra optimizations were performed. Therefore, we found the movement trajectory that maximized accuracy while keeping the probability of hitting the obstacle below a fixed limit. Using the optimized knots, we were also able to generate trajectories from the optimal motor command corrupted by signal-dependent noise and to compare these noisy trajectories to those observed (Sabes and Jordan 1997).

Two factors were varied in the simulations, the coefficient of variation of the signal-dependent noise (CVN) and the permitted collision probability (CP). Five levels of motor command noise with CVN of 0.15, 0.20, 0.22, 0.25, and 0.30 were examined with a CP of 2.3%. In addition, five CPs of 0.6, 1.3, 2.3, 3.6, and 6.6% were examined with the CVN of 0.22. Four further simulations were performed to determine the interaction of these factors, testing every combination of the extreme values of CP and CVN.

In accord with the empirical studies (Sabes and Jordan 1997), our analysis focused on the distribution of near-points, that is, the location on the path which is nearest to the obstacle tip. The angular deviation of the near-point from the obstacle axis, referred to as the near-point angle (δ), was examined as a function of obstacle angle (ϕ) (see Fig. 1*B*). Previously it has been shown (Sabes and Jordan 1997) that the near-point angle is linearly related to the deviation of the obstacle angle from a preferred axis (ω) with a slope (β); that is, $\delta = \beta(\omega - \phi + n\pi)$ for integer n, with n chosen such that $(\omega - \phi + n\pi)$ lies between $-\pi$ and π. In this cyclical linear model, $+\pi$ and $-\pi$ are equivalent, so for stability we fit the simulated data and the original data from Sabes and Jordan (1997) using the set of three lines $\delta = \beta(\omega - \phi + n\pi)$ for $n = -1, 0, 1$. For each possible value of ω, each data point was fitted to the line that gave the least-squared error. Standard linear regression was used to obtain the value of β associated with this ω, and the values of β and ω that gave the smallest total error were taken as the fitted values. The preferred axis ω can be interpreted as the axis that the near-points cluster toward, and the slope parameter β can be interpreted as the degree of clustering toward this axis, or the amount of skewing in the hand paths: If β were 0, the near-point distribution would be uniform and the movement paths symmetric, and higher values of β indicate greater clustering of near-points toward the preferred axis and more skew in the movement paths.

RESULTS

The simulations found the optimal path for each obstacle angle in a mean of 1070 cycles (SD 48.5). Overall, 0.5% of the movements performed by the simulations did not optimize as required and produced paths with too high a collision probability; these paths were excluded from further analysis. Typical optimal paths are shown in Fig. 2 (heavy dashed line) for eight obstacle angles from the simulation with a CVN of 0.22 and a CP of 2.3% (the middle setting of each parameter). The covariance ellipses along the path show two SDs from the mean, and, as required, do not intersect the obstacle, thus ensuring a collision probability at any point along the path below 0.023. The solid lines are the individual paths taken from the five subjects from Sabes and Jordan (1997). The simulated and actual trajectories are qualitatively similar and the near-points for the actual (filled circles) and simulation (empty circles) paths are also similar. For example, most obstacle angles (Fig. 2, *A–E*) show skewed simulated and observed paths in which the near-points fall to one side of the obstacle axis (fine dashed line), while others (Fig. 2, *F–H*) are more symmetric, with near-points aligned along the obstacle axis. It is not possible to compare the dispersion of the trajectories in more detail than by observation, because Sabes and Jordan's five subjects each performed only one movement at each obstacle angle.

A typical subject's observed distribution of near-points (circles and crosses) show a preferred axis (dashed line) for the two positions in the workspace (Fig. 3, *A* and *C*). The equivalent near-points for a set of noisy paths simulated from the optimal path show a similar preferred axis (Fig. 3, *B* and *D*). The near-point distribution for the optimal paths (not shown) is similar to that for noisy paths, but as expected the optimal near-points cluster closer together than the noisy near-point distribution. The preferred axis can also be seen in the plots of near-point angle against obstacle angle (*lower panels*, Fig. 3). The pattern of near-points was fit with a cyclical linear model (dashed lines) relating the deviation of the near-points from uniformity to the difference between the obstacle axis and a fixed preferred axis. This shows a good qualitative fit between

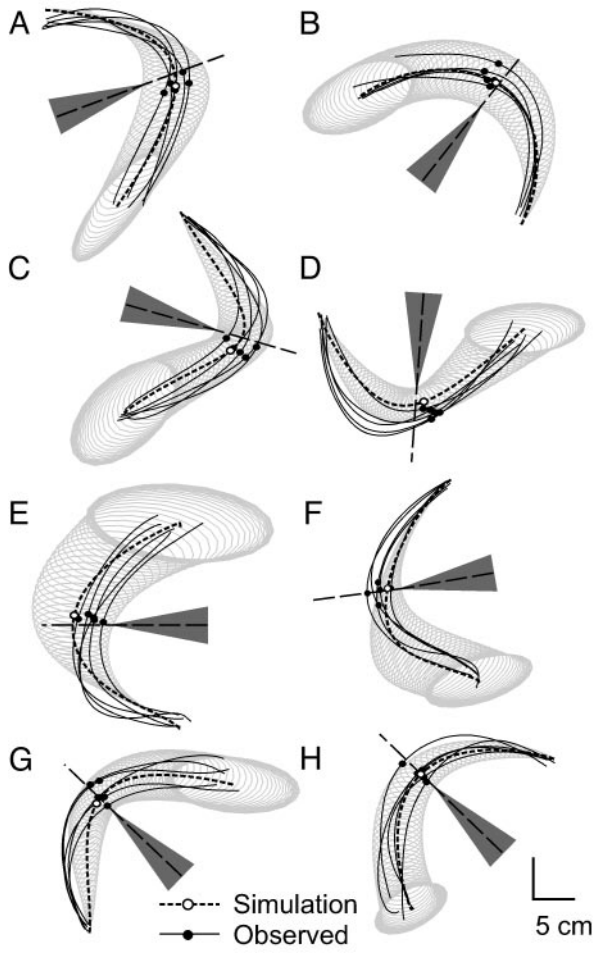

FIG. 2. Comparison of simulated and actual hand paths. The short dashed lines show the optimal path for each movement, and the pale ellipses show the 97.7% confidence ellipses of the path at each point. The open circles indicate the near-point of the optimal path. These data are taken from the simulation with a coefficient of variation of the signal-dependent noise (CVN) of 0.22 and a collision probability (CP) of 2.3%. Each thin black line shows a single trial from 1 of the 5 subjects performing the obstacle avoidance task (data from Sabes and Jordan 1997), and filled black circles show the near-points. Eight different obstacle angles are illustrated, all taken from position 1.

the cyclical linear models for this subject's near-point data and the simulated data.

In total, 13 simulations were carried out to compare the effects of changing the CVN or the CP. The performance of each simulation can be summarized in terms of the preferred axis and slope parameter from the cyclical-linear model, and the mean trace of the final covariance ellipse for the optimal trajectory. Figure 4 compares these performance measures for the nine simulations that hold one parameter at its middle setting and vary the other (solid lines and filled symbols), and for the four simulations examining combinations of the extreme values of each parameter (dashed lines and open symbols).

As expected, as the CVN increased or the CP decreased, the trace of the final covariance ellipse increased (Fig. 4, *E* and *F*). Similarly, the slope parameter increased with increasing CVN and with decreased CP (Fig. 4, *C* and *D*). Therefore, increasing the difficulty of the task either by increasing noise or by reducing collision probability increases the extent of path skewing as revealed by the slope parameter. The preferred axis was little affected by any change in the parameters of the simulations (Fig. 4, *A* and *B*). The orientation of the final variance ellipse was also unaffected: the mean change in orientation between simulations was 4.9° and the maximum change was 20°. The trends for the simulations that used extreme parameter values (dashed lines) are similar to those for the intermediate values (solid lines), and a significant interaction between the CP and CVN was only found for the slope parameter in position 1. The similar fits obtained across a range of CVN and CP values suggest that the simulation results do not depend heavily on the precise choice of these free parameters.

Confidence limits on each fit were obtained by bootstrapping (Efron 1982) and used to compare the two positions in the workspace (Fig. 4, position 1: gray symbols; position 2: black symbols). This revealed that the preferred axis was significantly higher in position two than position one for every simulation and that the slope parameter was significantly higher in position two than position one for some simulations.

To evaluate the performance of the simulations against the observed data, the slope parameter and preferred axis for the observed data (averaged over 5 subjects) were compared with the slope and preferred axis found for five near-point distributions (Fig. 5). These distributions were obtained from movements executed with noise using the simulation with a CP of 2.3% and a CVN of 0.22 (the middle setting of each parameter). *t*-tests revealed no difference between the preferred axes in position one ($t = 1.19$, df $= 8$, $P = 0.27$) or position two ($t = 0.92$, df $= 8$, $P = 0.39$), and no difference in the slope in

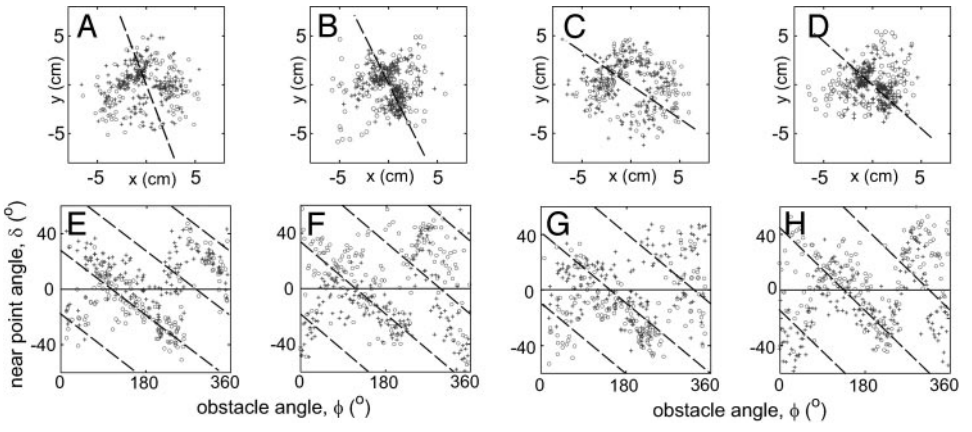

FIG. 3. Near-point distributions from observed and simulated trajectories in two positions (crosses are for clockwise movements and circles for counter-clockwise movements). *A*: near-point distributions in Cartesian space observed in human performance for position 1 (data from Sabes and Jordan 1997). Dashed lines show the preferred axis. *B*: near-point distributions on paths generated with noise by the simulation with CVN = 0.22 and CP = 2.3% in position 1. *C*: observed near-point distribution for position 2. *D*: simulated near-point distribution for position 2. *E–H*: the same data plotted as the near-point angle against the obstacle angle. The dashed lines show the fit to the cyclical linear model for each data set. The preferred axis is the intercept of the fit line with the *x*-axis and the slope parameter is the slope of the fit line.

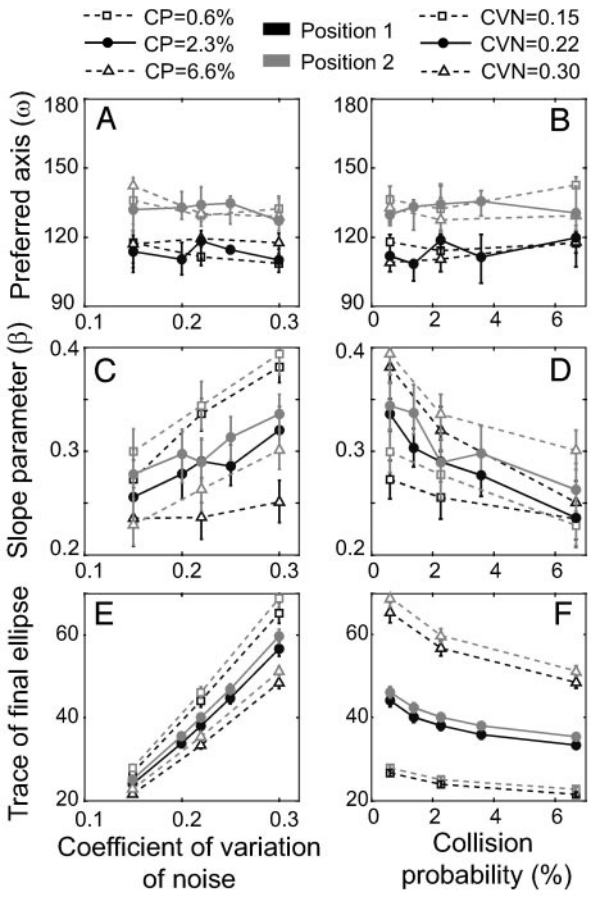

FIG. 4. Comparison of performance of the different simulations. Gray lines and symbols show data from position 1, and black lines and symbols show data from position 2. Simulation parameter settings are indicated in the legend and *x*-axis labels. *A* and *B*: preferred axis found for each simulation; error bars are 95% confidence limits. *C* and *D*: slope parameter for each simulation; error bars are 95% confidence limits. *E* and *F*: mean trace of the final variance ellipse for each simulation; error bars are SE.

position one ($t = 0.96$, df $= 8$, $P = 0.36$) or position two ($t = 0.49$, df $= 8$, $P = 0.64$).

The fit between the observed data and the mobility ellipse model proposed by Sabes and Jordan (1997) was also evaluated, and it was found that there was a significant difference between the observed and predicted preferred axes for position one ($t = 3.88$, df $= 4$, $P = 0.017$), but not for position two ($t = 2.66$, df $= 4$, $P = 0.056$).

DISCUSSION

We have modeled obstacle avoidance within the TOPS framework, in which the optimal trajectory is defined as the one that minimizes the mean-squared error at the end of the movement while keeping the probability of collision at each point along the path below a fixed limit. The optimal paths showed systematic deviations of the near-point (the point where the path came nearest to the obstacle) in accord with the empirical data. This deviation was characterized for both optimal and observed paths by the preferred axis, where the deviation is minimal, and a slope parameter that determines the extent of the deviation at other orientations. These parameters were not significantly different for the observed (Sabes and Jordan 1997) and optimal near-point distributions. In particular the optimal preferred axes were closer to the observed values when compared with the previous proposed model based on the mobility ellipse. Furthermore, the TOPS framework allows the central tendencies and its dispersion to be modeled rather than just the preferred axis. Therefore, the slope parameter can be obtained for the optimal trajectory and was shown not to be significantly different to that observed.

Assumptions

To determine the optimal trajectory with the TOPS model, we have made several assumptions. The first assumption is that the optimal feedforward trajectory is a good representation of the optimal feedback trajectory. We have modeled TOPS assuming no feedback, as feedback control with signal-dependent noise is currently an active but unsolved area of control theory (Lu and Skelton 2000). Interestingly, for linear systems the average optimal feedforward path is the same as the aver-

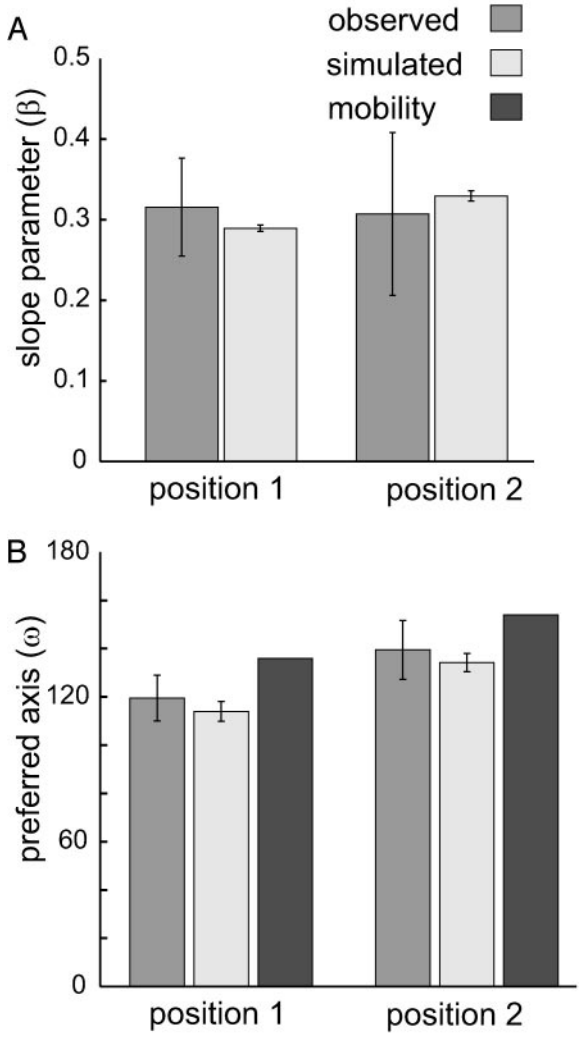

FIG. 5. Comparison of preferred axis and slope parameter between the simulated and observed data. *A*: mean slope parameter obtained for 5 noisy runs of the simulation with a CVN of 0.22 and a CP of 2.3% and for all 5 observed subjects. Error bars are 1 SD. *B*: preferred axis obtained from the same data with error bars as before, and the preferred axis predicted by the mobility ellipse.

age optimal feedback path. However, other characteristics of the path such as frequency may differ and the movements with feedback will be more accurate; this may explain why the movements observed by Sabes and Jordan have a smaller dispersion than the simulated movements (see Fig. 1). We have previously shown that TOPS provides a good model for feedforward movement, such as saccadic eye movements and fast arm movement (Harris and Wolpert 1998). We have also shown that we can model slower feedback movement and characterize the average trajectory, such as slow arm movement and drawing movements, in which TOPS reproduces the two-thirds power law (Harris and Wolpert 1998; Viviani and Terzuolo 1982). Therefore, it seems that even with the nonlinearity of the arm we can model the average behavior under feedback using a feedforward optimal control model.

The second assumption is in the specific parameters we have used to model the arm and the noise. We have previously shown that optimal trajectories are not sensitive to the exact parameters of the arm model (Harris and Wolpert 1998). Also, both the preferred axis and the slope parameter did not vary substantially over a wide range of collision probabilities or coefficient of variation of the noise. Therefore, our simulations are not heavily dependent on the choice of parameters.

Relation to other models

There is clearly a relation between the proposal of Sabes et al. (1997, 1998) and the TOPS model of obstacle avoidance, in that mobility and the consequences of signal-dependent noise are related. Given noise in the torque, the mobility ellipse describes the variability in the Cartesian acceleration of the hand, which Sabes and Jordan (1997) relate to the preferred axis. However, the mobility ellipse cannot predict the slope parameter because it only defines the axis that near-points cluster toward, not the degree of clustering. Although the ratio of the major-to-minor axis of the mobility ellipse is likely to be related to the slope (for example, a higher ratio will go with a higher slope), knowing the ratio does not allow us to predict the actual value of the slope parameter. In the TOPS model the precise way in which the noise plays through the motor system is determined not only by inertia, as in the mobility ellipse model, but also by the dynamics of the muscle and the moving arm. Because TOPS is able to model the whole trajectory of the obstacle avoidance movement, it is able to make more detailed and general predictions than the mobility ellipse model.

The TOPS model also differs from other models of trajectory planning (Flash and Hogan 1985; Uno et al. 1989), which describe only the optimal trajectory, and do not consider the dispersion of trajectories which result from performing movements in a noisy system. These previous models predict a path that passes as close as possible to the obstacle, which is not seen in studies of obstacle avoidance (Abend et al. 1982; Sabes and Jordan 1997). Although it would be possible to more accurately simulate curved movements using these models by placing via-points (Abend et al. 1982; Viviani and Terzuolo 1982), no model has been proposed which could determine their location. In contrast, we have shown that the TOPS framework allows for curved movements without any explicit specification of a via-point.

For TOPS to be an account of motor planning, several processes have to take place within the CNS. First, there must be signal-dependent noise on the motor command. As already mentioned, such signal-dependent noise can be seen behaviorally in isometric tasks (Schmidt et al. 1979; Slifkin and Newell 1999) and in movement tasks (Fitts 1966). At the neurophysiological level, several studies (for example, Andreassen and Rosenfalck 1980; Clamann 1969; Matthews 1996) have shown variability in firing of motor neurons, that is, a constant coefficient of variation in motor unit interspike intervals. Such variability together with the recruitment pattern can be shown, through simulation, to lead to signal-dependent noise in muscle force output (A. Hamilton, K. E. Jones, and D. M. Wolpert, personal communication). This arises from the following two sources: variability in the motor neuronal firing and the recruitment pattern. As force output increases, the size of units recruited increases, leading to larger variability.

Second, the cost of a movement must be evaluated. In the TOPS framework the cost is movement error, which is behaviorally relevant and simple for the CNS to measure, unlike the cost functions in many other optimal control frameworks. Although we have minimized the mean-squared error in a batch fashion over multiple trajectories, the CNS gets to experience one movement at a time. It would be possible to calculate the average error, but it is more likely that the error on each trajectory is used to update the controller, with a slow learning rate. This effectively averages over recently experienced movements without the need to store previous errors. There is evidence that the cerebellum represents the error signal of each movement. For example, Kitazawa et al. (1998) recorded from the cerebellum as monkeys reached to a target. They found that complex spike firing conveyed information about the error in hand position and proposed that this signal may be used for learning movements.

Third, the trajectory must be updated to reduce the cost. Although we have solved the optimization problem using a feedforward motor command, we do not believe that this is the mechanism used neurophysiologically to generate movements. In general, a feedforward optimal control problem can be reformulated as an optimal feedback controller. For example, Hoff and Arbib (1993) showed that a feedback controller can be constructed to generate minimum jerk trajectories and that this feedback controller can deal with perturbation during the movement, such as target movements. Therefore, the error signal into the TOPS model could be used to tune a feedback controller so that it generates more optimal solutions. This obviates the need to re-solve the optimal control problem each time a new movement is made and can also be used on-line to compensate partially for noise arising during the movement. Although the neural location of such adaptation is not known, we speculate that the cerebellum may combine both adaptive feedforward and feedback controllers which are updated by the error signals already described.

Finally, we believe that uncertainty is a fundamental problem in sensorimotor control. The motor system has to cope with incomplete information about the world and with noise in both its sensory inputs and its motor commands. Analysis of the role of uncertainty in the CNS has been valuable in understanding perceptual judgment (Britten et al. 1993; Gold and Shadlen 2000) and decision-making (Platt and Glimcher 1999). We suggest that examination of the role of noise and uncertainty in relation to the motor systems will form an important and productive theme in the future. We also believe that the

TOPS framework can be extended to almost any task. For example, to achieve accurate throwing, we could optimize the final position of the thrown object. This might favor certain types of correlation between the position and velocity of the hand at the moment of release, perhaps a faster speed if we release too early or a slower speed if we release too late in the trajectory. Therefore, by altering the motor command we could control aspects of the full covariance of the state of our body.

In conclusion, we have applied the TOPS framework to obstacle avoidance and have demonstrated that the model can accurately predict the pattern of behavior observed. We suggest that many types of movement can be planned by considering the statistics of action and that TOPS provides a potential unifying framework for understanding goal-directed action.

APPENDIX

Given an n-dimensional variable $\mathbf{x}$, with mean $\bar{\mathbf{x}}$ and covariance $\mathbf{P}_{xx}$, the unscented filter (Julier and Uhlmann 1995, 1996) can be used to estimate the distribution of $\mathbf{y} = f(\mathbf{x})$. To do this, we construct a set of $2n + 1$ points X, each of which is assigned a weighting W_i

$$X_0 = \bar{\mathbf{x}} \qquad W_0 = \kappa/(n + \kappa)$$
$$X_i = \bar{\mathbf{x}} + (\sqrt{(n + \kappa)\mathbf{P}_{xx}})_i \qquad W_i = 1/2(n + \kappa)$$
$$X_{i+n} = \bar{\mathbf{x}} - (\sqrt{(n + \kappa)\mathbf{P}_{xx}})_i \qquad W_{i+n} = 1/2(n + \kappa)$$

where $(\sqrt{(n + \kappa)\mathbf{P}_{xx}})_i$ is the ith row of the matrix square root of $(n + \kappa)\mathbf{P}_{xx}$. These points are chosen so that they give an unbiased representation of the mean and covariance of the original distribution. Each point is transformed through the nonlinear function f to give $\mathbf{y}_i = f(\mathbf{x}_i)$. The new mean $\bar{\mathbf{y}}$ and covariance $\mathbf{P}_{yy}$ can be found as

$$\bar{\mathbf{y}} = \sum_{i=0}^{2n} W_i \mathbf{y}_i \qquad \mathbf{P}_{yy} = \sum_{i=0}^{2n} W_i \{\mathbf{y}_i - \bar{\mathbf{y}}\}\{\mathbf{y}_i - \bar{\mathbf{y}}\}^{\mathbf{T}}$$

The new mean and covariance can be used to define a new set of points X, and the process is iterated for every time step.

In the case of the arm model, the state variable $\mathbf{x}$ has 12 dimensions: $\mathbf{x} = (q_1, q_2, \dot{q}_1, \dot{q}_2, \ddot{q}_1, \ddot{q}_2, tq_1, tq_2, t\dot{q}_1, t\dot{q}_2, u_1, u_2)$, where q_i is the joint angle, tq_i is the torque, and u_i is the motor command at the ith joint. The nonlinear function f was the forward model of the muscles and dynamics of the arm. We assume that the distribution of $\mathbf{x}(\kappa)$ is Gaussian, and so, following Julier and Uhlmann (1995), chose $(n + \kappa) = 3$. At the start of movement, the covariance of the state is a matrix of zeros. Variability in the state arises because at each time step the SD of the motor commands is set to be linearly related to the absolute value of the motor commands, with the constant of proportionality given by variable CV.

We thank P. Sabes for allowing us to use his published data and P. Vetter and R. van Beers for comments on the manuscript.

A. Hamilton is funded by a Brain Research Trust studentship. This work was funded by the Brain Research Trust, the Wellcome Trust, and the Human Frontiers Science Program.

REFERENCES

ABEND W, BIZZI E, AND MORASSO P. Human arm trajectory formation. *Brain* 105: 331–348, 1982.

ANDREASSEN S AND ROSENFALCK A. Regulation of the firing pattern of single motor units. *J Neurol Neurosurg Psychiatry* 43: 897–906, 1980.

BRITTEN KH, SHADLEN MN, NEWSOME WT, AND MOVSHON JA. Responses of neurons in macaque MT to stochastic motion signals. *Vis Neurosci* 10: 1157–1169, 1993.

CLAMANN HP. Statistical analysis of motor unit firing patterns in a human skeletal muscle. *Biophys J* 9: 1233–1251, 1969.

EFRON B. *The Jackknife, the Bootstrap, and Other Resampling Plans*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 1982.

FITTS PM. Cognitive aspects of information processing. 3. Set for speed versus accuracy. *J Exp Psychol* 71: 849–857, 1966.

FLASH T AND HOGAN N. The coordination of arm movements: an experimentally confirmed mathematical model. *J Neuroscience* 5: 1688–1703, 1985.

GOLD JI AND SHADLEN MN. Representation of a perceptual decision in developing oculomotor commands. *Nature* 404: 390–394, 2000.

HARRIS CM AND WOLPERT DM. Signal-dependent noise determines motor planning. *Nature* 394: 780–784, 1998.

HOFF B AND ARBIB MA. Models of trajectory formation and temporal interaction of reach and grasp. *J Mot Behav* 25: 175–192, 1993.

HOGAN N. The mechanics of multi-joint posture and movement control. *Biol Cybern* 52: 315–331, 1985.

JULIER SJ AND UHLMANN JK. A new extension of the Kalman filter to nonlinear systems. In: *Proceedings of AeroSense: The 11th International Symposium on Aerospace/Defense Sensing, Simulation and Controls, Session: Multi Sensor Fusion, Tracking and Resource Management II* [letter], Orlando, FL, 1995.

JULIER SJ AND UHLMANN JK. A general method for approximating nonlinear transformations of probability distributions [Internet Publication. Letter]. Robotics Research Group, Oxford University, 1996. http://www.robots.ox.ac.uk/~siju/work/publications/A4_size/Unscented.zip [2001 May 16]

KAWATO M. Bi-directional theory approach to integration. In: *Attention and Performance 16: Information Integration in Perception and Communication,* edited by Inui T and McClelland JL. Cambridge, MA: The MIT Press, 1996, vol. XVII, p. 335–367.

KIRKPATRICK S, GELATT CD, AND VECCHI MP. Optimization by simulated annealing. *Science* 220: 671–680, 1993.

KITAZAWA S, KIMURA T, AND YIN PB. Cerebellar complex spikes encode both destinations and errors in arm movements. *Nature* 392: 494–497, 1998.

LU J AND SKELTON RE. Robust variance control for systems with finite-signal-to-noise uncertainty. *Automatica* 36: 511–525, 2000.

MATTHEWS PB. Relationship of firing intervals of human motor units to the trajectory of post-spike after-hyperpolarization and synaptic noise. *J Physiol* 492(Pt. 2): 597–628, 1996.

METROPOLIS N, ROSENBLUTH A, ROSENBLUTH M, TELLER A, AND TELLER E. Equation of state calculations by fast computer machines. *J Chem Phys* 21: 1087–1092, 1953.

PLATT ML AND GLIMCHER PW. Neural correlates of decision variables in parietal cortex. *Nature* 400: 233–238, 1999.

PRESS WH, TEUKOLSKY SA, VETTERLING WT, AND FLANNERY BP. *Numerical Recipes in C*. Cambridge, UK: Cambridge Univ. Press, 1992.

SABES PN AND JORDAN MI. Obstacle avoidance and a perturbation sensitivity model for motor planning. *J Neuroscience* 17: 7119–7128, 1997.

SABES PN, JORDAN MI, AND WOLPERT DM. The role of inertial sensitivity in motor planning. *J Neuroscience* 18: 5948–5957, 1998.

SCHMIDT RA, ZELAZNIK H, HAWKINS B, FRANK JS, AND QUINN JTJ. Motor-output variability: a theory for the accuracy of rapid motor acts. *Psych Rev* 47: 415–451, 1979.

SLIFKIN AB AND NEWELL KM. Noise, information transmission, and force variability. *J Exp Psych Hum Percept Perform* 25: 837–851, 1999.

UNO Y, KAWATO M, AND SUZUKI R. Formation and control of optimal trajectory in human multijoint arm movement. Minimum torque-change model. *Biol Cybern* 61: 89–101, 1989.

VAN DER HELM FCT AND ROZENDAAL LA. Musculoskeletal systems with intrinsic and proprioceptive feedback. In: *Biomechanics and Neural Control of Posture and Movement,* edited by Winters JM and Crago PE. New York: Springer-Verlag, 2000, p. 164–177.

VIVIANI P AND TERZUOLO C. Trajectory determines movement dynamics. *Neuroscience* 7: 431–437, 1982.