

Numerically Estimating Internal Models of Dynamic Virtual Objects

G. ROBLES-DE-LA-TORRE

International Society for Haptics

and

R. SEKULER

Brandeis University

Precise manipulation of objects is ordinarily limited by visual, kinesthetic, motor, and cognitive factors. Specially designed virtual objects and tasks minimize such limitations, making it possible to isolate and estimate the internal model that guides subjects' performance. Subjects manipulated a computer-generated virtual object (vO), attempting to align vO to a target whose position changed randomly every 10 s. To analyze the control actions subjects use while manipulating the vO , we benchmarked human performance against that of ideal performers (IPs), behavioral counterparts to ideal observers used in sensory research. These comparisons showed that subjects performed as feed-forward, predictive controllers. Simulations with degraded-IPs suggest that human asymptotic performance was not limited by imprecisions of vision or of motor timing, but resulted mainly from inaccuracies in the internal models of vO dynamics.

Categories and Subject Descriptors: H.1.2 [Information Systems]: Models and Principles—*User/machine systems*; J.3 [Computer Applications]: Life and Medical Sciences—*Biology and genetics*; I.5.4 [Computing Methodologies]: Pattern Recognition—*Applications*

General Terms: Human Factors, Performance, Experimentation

Additional Key Words and Phrases: Dynamics, human cognition, human information processing, ideal performer, internal model, virtual object, virtual reality

1. INTRODUCTION

Understanding how perceptual, motor, and cognitive processes contribute to action is a basic goal of neuroscience. This knowledge is also essential for the design and evaluation of virtual environments. Here, we present a quantitative framework for studying cognitive processes in action, specifically how humans acquire an internal model of a dynamic virtual object. Our methodology minimizes limitations imposed by motor or perceptual systems, allowing a more direct connection between observed performance, on the one hand, and the information that guides subjects' actions, on the other. To anticipate, benchmarking human performance against a computationally well-defined theoretical strategy makes it possible to characterize human performers as feed-forward controllers that seek to minimize their effort. This in turn allows a numerical estimation of the internal model humans use when manipulating an object, and of experience-driven changes in that model. The result is a quantitative account of cognitive and neural processes expressed in the language of dynamics [Ogata 1978].

Authors' address: G. Robles-De-La-Torre, Department of Psychology, Queen's University, Kingston, ON, Canada K7L 3N6; email: Gabriel@RoblesDeLaTorre.com; R. Sekuler, Volen Center for Complex Systems, Brandeis University, Waltham, MA 02454; email: vision@brandeis.edu.

Permission to make digital/hard copy of part of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of publication, and its date of appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or fee.

© 2004 ACM 0000-0000/04/1000-0102 \$5.00

To characterize how human performers exploit internal models to plan and generate appropriate actions [Wolpert and Ghahramani 2000; Wolpert and Flanagan 2001; Davidson and Wolpert 2003; Flanagan et al. 2003] when guiding objects under novel conditions [Atkeson 1989; Davidson and Wolpert 2003], we asked human subjects to manipulate a computer-animated virtual object. This virtual object (vO) was a high contrast, black horizontal bar, presented on an uniform, white background. The vO extended rightward from the computer display's left side. Using a pair of buttons, a subject tried to bring the leading edge of this bar into alignment with a single vertical target line that lay slightly above the vO 's path. Pressing and holding one button drove the object's tip slowly rightward; pressing and holding the other button drove the object's tip leftward at the same rate. Every 10 s, the horizontal position of the virtual target's tip changed randomly, which also randomized the horizontal distance separating the virtual object from the target. This random variation made it impossible for subjects to perform well by merely learning how long to hold a button down, or by learning some fixed sequence of button presses.

2. MATERIALS AND METHODS

Subjects pressed either one or the two control buttons as frequently, and held a button down as long, as was needed to bring the virtual object and target into alignment. Subjects were instructed to do this as quickly and accurately as possible, but were not instructed about the strategy they ought to adopt. Subjects were free to use as many button presses as needed to perform the task. We allowed this freedom, even though it would complicate data analysis, in order to avoid biasing subjects into adopting any particular strategy. We wanted to observe how subjects' strategies evolved as free of constraints as possible. We then sought to characterize those strategies, using system analysis and computer simulation.

Subjects' motor responses were limited to button presses in order to minimize complications associated with learning, planning, and executing complex motor actions. Together with the stimulus' highly discriminable visual features, the simplification of the task's motor demands was intended to facilitate links between empirical performance and the subject's learning process.

An ideal observer is an algorithm that produces optimal detection or discrimination in a well-specified sensory task. Such algorithms make it possible to define sensory tasks with high precision and also to identify factors that limit human performance [Geisler 1989; Altes 1989; Barlow 1980]. We used a behavioral counterpart to an ideal observer, an ideal performer (IP), which is a quantitative benchmark against which subjects' performance and improvement with training could be assessed. Recently, a Kalman filter was used as an IP to assess human behavior in a visuomotor task. In that work [Baddeley et al. 2003], the Kalman filter represented the optimal performer in least-squares/maximum-likelihood sense. The filter used the errors on previous trials to estimate the relevant parameters on the current trial.

In contrast, we are interested in examining human performance in a framework that is closer to dynamics than to statistics. It is then necessary to have analysis tools to examine human interaction with systems whose state changes over small time scales. For this, we designed IPs that achieve optimum performance for a given control strategy within a single trial. Implemented in Matlab, the IPs are predictive control systems, which incorporate complete knowledge of the virtual object's dynamics, perfect information about the object's position and velocity, and about the current position of the target. Moreover, an IP's perfect time base and memory allows it to push the appropriate control button at exactly the right instant, hold it for the right duration, and release it at just the right time. To test various hypotheses about human subjects' strategies, we devised alternative IPs, each of which pursued a different, specific goal as it performed our task.

By assigning a virtual mass (m) and a virtual viscous resistance (b) to the computer-animated object, the computer code defined the object's dynamics, that is, how it would respond to virtual forces. In our experiment, $m = 14.17$ units and $b = 17$ force units sec/mm. With these parameters, if a subject pressed and held down one response button for an entire 10-s trial, the object accelerated to 7.7 deg/s by the trial's end. Note that even this highest attainable speed fell within the range where changes in velocity are discriminated optimally [McKee and Watamaniuk 1994]. Moreover, the small vertical separation between vO 's tip and target line was meant to insure excellent discrimination of object position relative to the target line [Waugh and Levi 1995].

2.1 Virtual Object, Targets, and Control Signal

The stimuli were presented at a viewing distance of 0.5 m. The virtual object was a horizontal, black bar, 19.1' high, presented on a uniform, white background. The luminances of bar and background were 1.2 and 68.7 cd/m², respectively. Distances between target and object were drawn from a uniform distribution, ranging from 0.4 to 14.5 deg visual angle. The target line (55' long) was always displayed in the area above the virtual object. When the virtual object and the target were aligned, the vertical separation between them was 22'. Movements of the vO 's tip were governed by the differential equation:

$$m \frac{d^2x(t)}{dt^2} = C_S(t) - b \frac{dx(t)}{dt} \quad (1)$$

where m and b are the virtual object's mass and viscous resistance, respectively; $x(t)$ is the position of the object's tip at time t ; and $C_S(t)$ is the subject's control signal, the virtual forces generated by the two control buttons. In our experiment, the control signal $C_S(t)$ in Eq. (1) comprised a ramp function of the form $F(t) = \pm 250(t - t_i)$ force units, where t_i is the time of onset of a button press, in seconds. By convention, $F(t)$ is negative when the subject presses the left button, and positive if the subject presses the right button. Because releasing the button reset $F(t)$ to zero, each button press applied a new ramp function to the object. Regarding the selection of a dynamics with substantial viscosity, we considered that using a simple, energy-dissipative dynamics would help us understand the feasibility of the approach we propose, while resulting in an experimental situation that was rich enough to be meaningful.

The virtual object's dynamics can be represented also by $h(t)$:

$$h(t) = \frac{1 - e^{-bt/m}}{b} \quad (2)$$

which is the impulse response function that results from solving Eq. (1) with $C_S(t) = \delta(t)$ (Dirac's delta) and initial conditions equal to zero.

The position of the virtual object was computed from Eq. (1). This equation was integrated in real time using a fourth-order Runge–Kutta method. Both the position of the virtual object and the subject's control signal, $C_S(t)$, were recorded for off-line analysis. Data were acquired at 1 kHz using a special purpose computer interface.

3. EXPERIMENT

We tested 19 subjects, 9 females and 10 males, 18–30 years old; all were naive about the experiment's purpose, were paid for their participation, and gave informed consent prior to testing. Subjects were instructed to press the control buttons as often and as long as needed to bring the virtual object and target into alignment, as quickly and as accurately as possible. A session started with six practice trials after which subjects received a total of 240 experimental trials. Trials were distributed over four 60-trial blocks, separated by three short breaks. As indicated before, the virtual object's dynamics followed Eq. (1), with $m = 14.17$ mass units, $b = 17$ force units sec/mm.

4. DATA ANALYSIS

4.1 Effort-Minimizing IP

The term “effort” refers here to the number of control actions that are used to align the object to the target. An “effort-minimizing” IP minimizes its effort by generating a single, sustained button press that accelerates the object to the desired velocity. Then the IP releases the button, allowing the object to coast smoothly into alignment with the target. The νO 's velocity then decreases exponentially, approaching zero at alignment. It is important to note that IP's button release time varies with the distance that the virtual object must travel. The optimal button release time for the νO 's we used is given by

$$t_R = \sqrt{\frac{2Db}{250}} \quad (3)$$

D is the distance to the target.

4.2 Time-Minimizing IP

This class of IP minimizes the time needed to align object and target. It accelerates νO to its maximum, and then alternates left and right button presses to brake νO optimally.

4.3 Measure of Performance

The trialwise error measure

$$E = \sum_{i=1}^n \frac{1 + |x_S(t_i) - D|}{1 + |x_{IP}(t_i) - D|} \quad (4)$$

was used to evaluate subjects' performance improvement in each trial of the experiment. This measure compares the position of the subject-controlled νO to the position of the same object under IP control. D is the distance to the target, $x_S(t_i)$ and $x_{IP}(t_i)$ are the positions at time t_i of the νO 's tip when under a subject's and the IP's control, respectively. n is the number of selected points to be sampled during the trial. The summation is evaluated over the duration of the trial, starting at the first button press.

4.4 Dynamics-Normalized Performance Measure

Eq. (4) compares the νO 's trajectory under the control of subjects to its trajectory under IP's control. This is an adequate overall measure, but failed to capture some important information about subjects' behavior. For example, if a subject made very short presses of the buttons, the νO 's damping and inertia would keep the νO 's trajectory from reflecting those button presses. We solved this problem by designing a performance measure that directly compared subjects' control signals with the control signals of IP's. The derivation of this dynamics-normalized measure is presented in Appendix A. Figure 1 gives a schematic, intuitive description of this measure.

This dynamics-normalized measure affords several advantages. In the example shown in Figure 1, we compared subject's performance to the performance of effort-minimizing IPs. But the normalized measure can be used to compare a subject's performance to that of any IP. Also, regardless of the target's distance, the IP's dynamics-normalized performance is always equal to the νO 's impulse response function. As a result, subjects' performance can be compared to the IP's by using always the same reference. Much of our data analysis exploited the dynamics-normalized measure to compare subjects' behavior to the effort-minimizing IP's only. In what follows, “normalized performance” refers to this dynamics-normalized measure of performance.

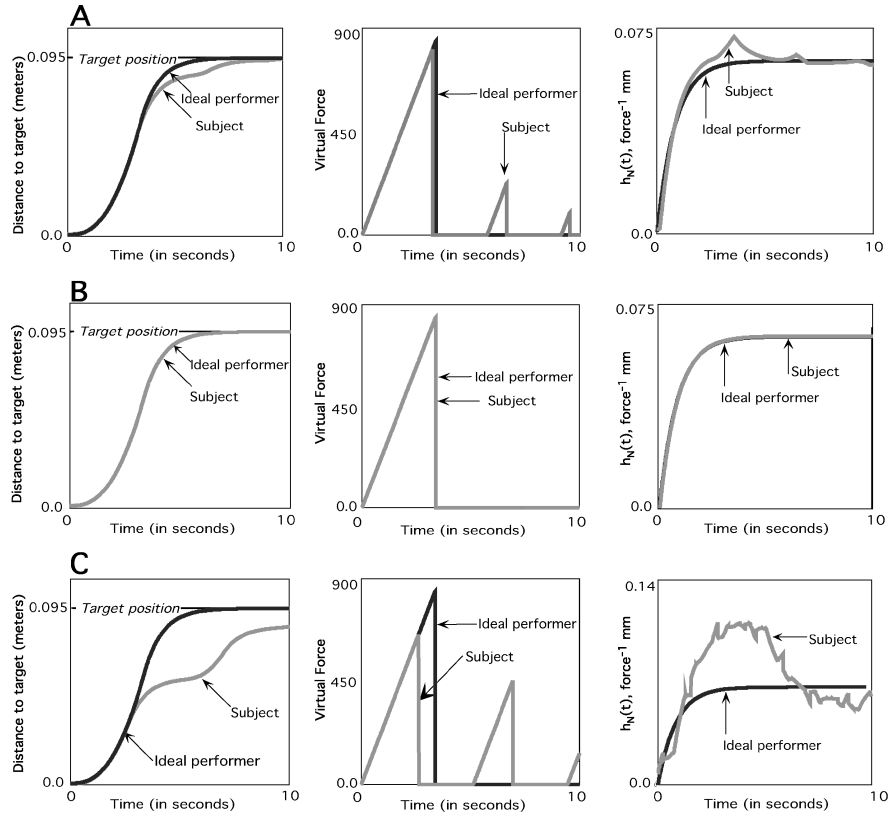


Fig. 1. Schematic explanation of dynamics-normalized performance. (A) Left panel: the trajectory of the virtual object when a simulated subject aligns the object to a target (gray curve). Compare it to the trajectory of the object when an effort-minimizing IP aligns the object to the same target (black curve). Central panel: the control actions that the subject (gray curve) and the IP (black curve) used to align object and target. Note that the IP used a single button press, while the simulated subject used three. Right panel: subject's and IP's control signals (central panel) were transformed into dynamics-normalized form using the algorithm shown in Appendix A. The IP's normalized performance (black curve) is the same as the virtual object's impulse response function (Eq. 2). Because subject's control signal (central panel) differs from the IP's, the subject's normalized performance deviates from the virtual object's impulse response function. (B) Left panel: both the subject and the IP use the same control signal (central panel) to align the object to the target. As a result, their dynamics-normalized performances are identical (right panel). (C) Left panel: the subject aligns the object to the target using a control signal (central panel) that differs considerably from the IP's. This yields the subject's dynamics-normalized performance (right panel), which is very different from the IP's. Note that the IP's dynamics-normalized performance is always equal to the virtual object's impulse response function. This equality is independent of the distance to the target.

4.5 Outliers and Calculation of Trialwise Errors

When applying the performance error measures described above for each subject, a mean trial error μ and its standard deviation σ were calculated. Trials in which performance errors were outside the interval $\mu \pm 2\sigma$ were eliminated. Performance measure errors were then cumulated across trials.

The resulting error curve was fitted with two straight lines. Each line estimates the error in trials spanned by the line. The first line fits the error from trial₁ to trial_i, and the second does it from trial_{i+1} to the last trial. Trial_{i+1} defines the point in which asymptotic performance is achieved [Bogartz 1971]. The curve-fitting procedure is set to automatically minimize error given trial_i and the related slopes.

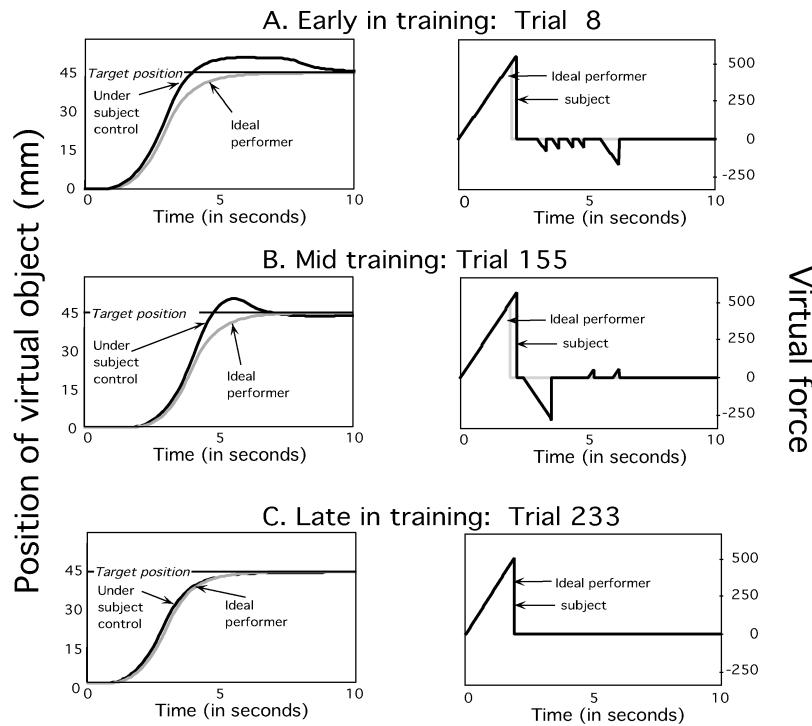


Fig. 2. Samples of a typical subject's performance at various times during training: (A) Early in training; (B) in the middle of training; and (C) late in training. In each case, the trial began with the virtual object approximately 4.8 degrees visual angle from the target (represented by the horizontal line). The left panels show the trajectory of the vO under subject's and IP's control. The right panels show the corresponding control signals that generate these trajectories. With practice, the human subject's behavior approximates that of the effort-minimizing IP.

Table I. Performance Improvement with Practice Measured Relative to Each IP

Ideal Performer	Mean Initial Trialwise Error	Mean Asymptotic Trialwise Error	Mean Number of Trials to Asymptote
Time-minimizing	136.4 ± 26	95.1 ± 8	99.0 ± 9.37
Effort-minimizing	74.6 ± 8	55.5 ± 5	115.4 ± 9.65

Mean \pm one standard error.

The slope of the first line defines a quantity that we call the initial trialwise error; the slope of the second line defines the asymptotic trialwise error.

5. RESULTS

5.1 Basic Behavioral Results

With training, subjects' performance improved relative to that of an IP. Figure 2 shows samples of a typical subject's performance at three different stages of training. Early on, as subjects tried to guide the virtual object into alignment with the target they tended to overshoot or undershoot the target position. If sufficient time remained during the 10-s trial, subjects made additional button presses to correct errors. With training, subjects, including the one whose data are shown in Figure 2, greatly reduce the number of control actions (button presses) used to align object and target. Subjects approach but do not reach the performance of the feed-forward IP. Table I shows mean trialwise errors, computed using Eq. (4).

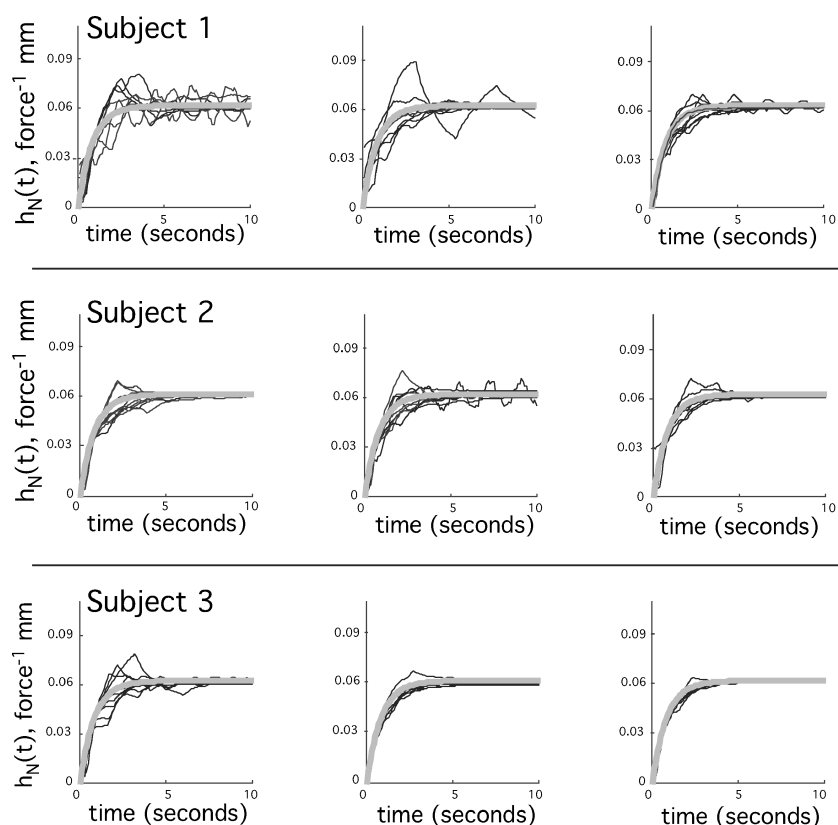


Fig. 3. The dynamics-normalized performance measure shown at different stages of training, for three typical subjects. Distances to target were greater than 1.61 deg visual angle. In each panel, the virtual object's impulse response function (Eq. 2) is shown by a thick, gray curve, and the normalized performance measure by thinner lines, one for each of ten consecutive trials. Each row plots the performance of one subject early in training (left panel), in mid-training (central panel) and late in training (right panel). As training progresses, dynamics-normalized performance approaches the virtual object's impulse response function.

As the data in Table I suggest, with practice, subjects' performance approached both IP's performance, but came significantly closer to the effort-minimizing IP ($p < 0.001$). Every subject showed this advantage of the effort-minimizing IP over its time-minimizing counterpart. This strongly suggests that subjects were using a strategy that more closely approximated the strategy of the effort-minimizing IP.

But exactly how closely did human subjects develop this strategy? We used the dynamics-normalized error measure to gain further insight into this (see Section 4.4). This measure highlighted several important aspects of subjects' behavior. When the initial displacement of object and target was between 0.40 (our smallest starting distance) and 1.61 deg visual angle, the mean error per trial was approximately three times the error for larger starting distances. For displacements below 1.61 deg, mean initial and asymptotic trialwise errors of all subjects were 131 and 107, respectively. In contrast, for displacements above 1.61 deg mean initial and asymptotic trialwise errors were 44 and 34. This difference was consistent for all subjects.

Displacements greater than 1.61 deg visual angle comprised 80% of all trials. For these large displacements, Figure 3, shows the normalized performance, $h_N(t)$, for three representative subjects at various points in training. For comparison, these figures show the object's impulse response function, $h(t)$ (Eq. (2)).

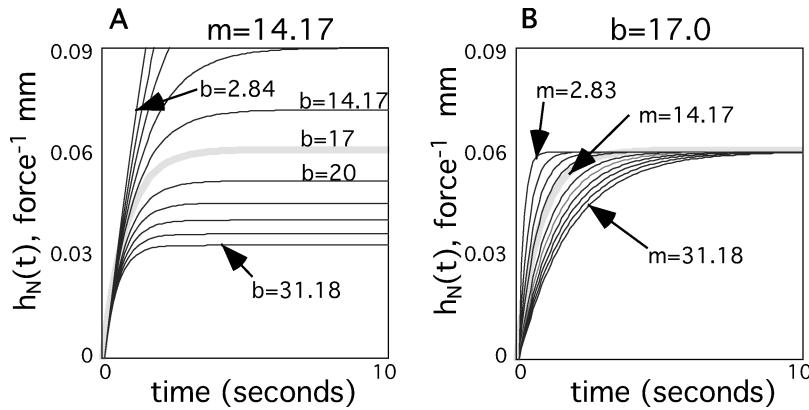


Fig. 4. Impulse response functions for different virtual objects. (A) When the mass stays constant ($m = 14.17$ here, value used for virtual object), variation of b changes the curves' asymptotic value, which is located at $\frac{1}{b}$. Note that even slight departures from $b = 17$ (gray, thick curve, value used for the virtual object) produces distinctly different curves. (B) When b , viscous resistance, is kept constant at 17 (value used in virtual object) and m is varied, the asymptotic value remains the same, but the curve approaches asymptote more rapidly with smaller mass and more slowly with larger mass.

With training, a subject's normalized performance approaches the vO 's impulse response function. For this to happen, the subject's control signal, C_S , must approach the control signal C_{IP} , its IP counterpart (see Section 4.4 for the reason why this is so). This finding reinforces the idea that with practice, subjects increasingly exhibit the feed-forward behavior of an effort-minimizing IP.

Figure 4 helps evaluate how well subjects' normalized performance approximates reference curves for objects with different dynamics. This figure shows the dynamics associated with various combinations of viscous resistance (b) and mass (m). These curves are also reference curves for assessing subjects' normalized performance, which is expressed in the same units (see Section 4.4). Panel A shows how the reference curves vary with viscous resistance, b , when $m = 14.17$, vO 's virtual mass. Note that curves' asymptotic value is $1/b$. Panel B shows how the reference curves vary with m , when $b = 17$, vO 's viscous resistance. Although all curves approach the same asymptote, curves for lighter masses reach that asymptote more quickly than curves for heavier masses. Both families of curves suggest that even small deviations from $m = 14.17$ and $b = 17$ would produce behavior distinctly different from what our subjects actually do. Compare subjects' performance (Figure 3) to the reference-normalized performance curves (Figure 4). These reference curves indicate that subjects' performance is a very good approximation to the feed-forward, effort-minimizing IP's. Given this, we propose that the subjects behave essentially as an effort-minimizing, feed-forward controller.

Figure 5 shows $h_N(t)$ for small starting displacements, below 1.61 deg, at different stages of training. The results are for the same three subjects represented in Figure 3 for larger starting displacements. Comparing both figures shows that subjects behave very differently in each case. The most obvious reason for this difference is that shorter travel distances afford more time for corrective actions. Because it took little time to align the vO and the target, subjects had enough time before the trial's end to execute many small, feedback corrections, which would produce the high-frequency components seen in this figure. Of course, here subjects would not be operating as feed-forward predictors. Another possibility, though, is suggested by the performance of Subject 3 in which the general shape of $h_N(t)$ is very similar to the impulse response function of the appropriate vO , but with a different limit value. This similarity indicates that this subject used a single button press, but misjudged the appropriate release time. This suggests that subjects had to sample the movement of the virtual object for some minimum time to

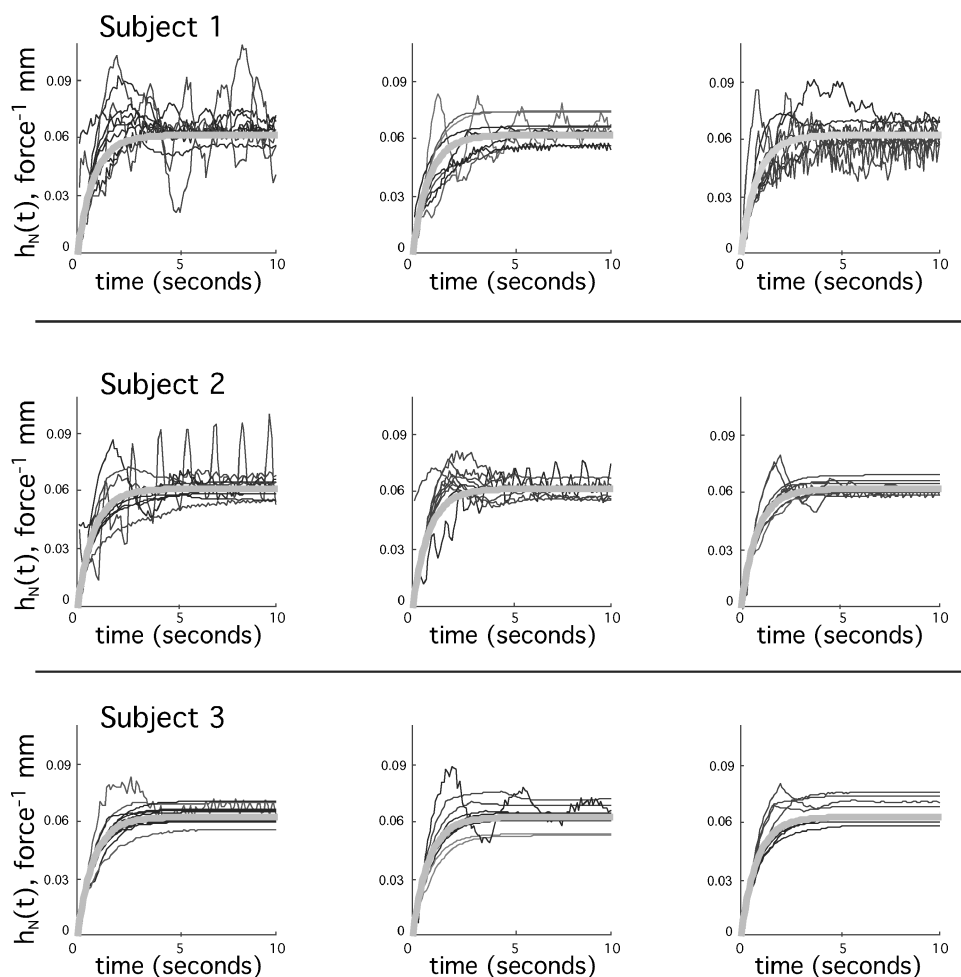


Fig. 5. The influence of small target distance (between 0.4 and 1.61 deg visual angle) on dynamics-normalized performance at different stages of training: early (left panels), mid-training (central panels) and late in training (right panels). Data are from the same subjects shown in Figure 3. Each row shows the performance of one subject. In each panel, the virtual object's impulse response function (Eq. 2) is shown by a thick, gray curve, and the normalized performance measure by thinner lines, one for each of ten consecutive trials. Compared to the corresponding panels in Figure 3, for small target distances the normalized performance measure differs substantially from the impulse response function, and it changes only weakly, if at all as training progresses.

gauge its velocity [Dzhafarov et al. 1993; Hohnsbein and Mateeff 1998], and/or that they needed a certain time to complete their motor planning, even after they initiated a button press.

Performance with shorter and longer starting separations of object and target differed in another way, too. For larger starting distances, training produced a significant decrease in the mean trialwise error ($p < 0.02$), but with the shorter starting distances, training's effect was considerably weaker ($p > 0.10$). This weaker effect is consistent with the idea that closed-loop corrections, based on visual feedback, play a larger role when starting distances are small than when they are large. Because we were interested mainly in practice-induced changes in performance, subsequent data analysis focused on trials whose starting distances were 1.61 deg visual angle or greater. These constituted 80% of all trials.




Fig. 6. (A) Initial (squares) and asymptotic (circles) mean trialwise normalized performance errors for 18 subjects. The shaded band in each column encompasses the mean trialwise error plus/minus one standard error. (B) Mean trialwise errors for IPs who mislocalize the target by varying amounts. Each data point is the mean of 1,000 simulated trials. The best-fit line ($r^2 = 0.98$) shows that with mislocalization of 30%, a noisy IP matches human subjects' mean initial trialwise error. At the end of the experiment, the subject's mean error corresponds to that of an IP who mislocalizes by 25%. Data for only 18 subjects are shown because one subject had atypically high errors.

Figure 6(A) shows the mean trialwise error near the beginning of the session (points in the left column) and at the end of the session (points in the right column). Although trialwise error declines with practice, even after 240 trials it has not reached zero. This residual error could come from various sources: noise in either the visual or motor aspect of the task, errors in motor planning, incomplete learning of the vO 's dynamics, or some combination of these factors. To examine contributions of visual and motor noise we used sensitivity analysis, examining the consequences of degrading particular components of an effort-minimizing IP. Specifically, then, we compared the degraded-IP's performance against that of our human subjects. Because of added noise, the effort minimizing IP necessarily generated more than one control action when aligning vO and target. These additional, corrective control actions were needed because the degraded-IP's first control action was no longer perfect. To decide when additional control actions were needed, the degraded-IP needed a criterion for determining when the decelerating vO actually ceased to move. We adopted a criterion based on human thresholds for changes in velocity [McKee and Watamaniuk 1994]. Specifically, motion was considered to have stopped when the vO 's velocity declined to 10% of its velocity when the button was released,

To clarify the possible role that might be played by defects in perception, we forced the IP to misjudge the distance to the target at time t , $D(t)$. One thousand trials were simulated for each of various levels of misperception. These perceptually impaired IPs were meant to mimic human subjects whose visual thresholds for localization were unusually elevated. Human judgments of changes in separation between two targets vary with the size of the starting separation between the targets. The Weber fraction for localization (WF_L) is a dimensionless value that expresses sensitivity to change in separation relative to initial separation, a quantity that in humans is about 3%. Results of the simulations defined a set of equivalencies between human subjects' performance, on the one hand, and the performance of a visually noisy, but otherwise IP, on the other. Figure 6B shows that subjects' mean asymptotic

error, 34.20, corresponds to the performance of an IP whose WF_L was 25%. This is about eight times the actual errors made by human subjects in comparable judgments [Burbeck and Yap 1990; Burbeck 1987, 1986]. To mimic the performance of our poorest subject (Figure 6A), the IP's localization would have to be more than ten times worse than would be expected from human subjects; to mimic even our best subjects would have to mislocalize to an extraordinary degree, more than six times normal. It is unlikely, therefore, that visual noise, in the form of misjudged distances, is the sole cause of human subjects' imperfect performance in our task.

In a corresponding set of simulations, the IP's vision was restored to perfection, but accuracy of its motor timing was degraded. On each simulated trial a random Gaussian variate ($\mu = 0$) was added to the moment at which the IP released the button. To match the asymptotic mean error rate of our human subjects the IP's release time had to be perturbed by Gaussian noise with $\sigma = 120$ ms. This value is about seven times the value expected from human subjects in such a task [Keele et al. 1985]. So, like visual noise, motor error is unlikely to be the sole explanation for our subjects' imperfect performance. But what happens when the two sources of noise operate at the same time?

To examine this possibility, we created IPs in which both visual and motor aspects were degraded, such IPs made localization errors and then made errors in release time. Various combinations of the two errors were able to match subjects' performance. To take one case, the mean error of human subjects would be matched by the impaired IP if the IP's WF_L were 15% and the noise in its timing were $\sigma = 70$ ms. Smaller visual errors required even greater motor error, and vice versa. All combinations that successfully reproduced subjects' performance required at least one error value that was implausibly large. The implausibility that human subjects' visual or motor systems would produce errors as large as those required by our simulations suggests that subjects' imperfect asymptotic performance comes largely from basic inaccuracies in subjects' planning of their motor action. These inaccuracies could involve either the information subjects acquire about the vO 's dynamics (i.e., subjects' internal model of the vO), or how this information is used to generate action. For example, a poor internal model would lead subjects to produce control actions that would inaccurately position the vO .

5.2 Dynamics-Normalized Performance and Subjects' Internal Model of the Virtual Object

To align the vO to the target, the IP uses its perfect knowledge of the vO 's dynamics. This knowledge is encapsulated in the object's impulse response function (Eq. (2)), and constitutes the IP's internal model of the virtual object (Appendix B). As explained in Section 4.4, the IP's dynamics-normalized performance is always equal to the object's impulse response function. This means that the IP's dynamics-normalized performance can be interpreted as an estimate of the IP's internal model of the virtual object. As we have seen, subjects approached the effort-minimizing IP's behavior as practice progressed. Then, it is reasonable to interpret a subject's normalized performance measure as an estimate of his/her internal model of the virtual object. We propose that the normalized performance measure is not only useful to assess learning, but that it can be used to estimate how similar the subjects' internal model is to the vO 's dynamics (see details in Appendix B).

6. DISCUSSION

We have examined human interactions with dynamically realistic virtual objects. By minimizing effects of motor and perceptual components, and by eliminating task-relevant kinesthetic and force-feedback information, it was possible to connect performance mainly to cognitive factors such as subjects' learning, motor planning, and internal models of the vO 's behavior.

Human subjects behaved asymptotically like effort-minimizing, feed-forward predictors of the object's responses. Presumably, operating in feed-forward mode requires an internal model that contains information about the virtual object's dynamics or some equivalent. The normalized performance measure

provides a way to estimate the parameters of such an internal model by comparing it to the dynamics of the virtual object.

In spite of being instructed to align the vO to the target as quickly and as accurately as possible, subjects' performance seemed to optimize effort, which is the number of actions made, rather than time. This suggests, among other possibilities, that subjects find it difficult to adopt a time-minimizing strategy. However, subjects may be reluctant to adopt a braking strategy such as the time-minimizing IP's. If this happened, subjects would apply virtual forces only in the direction of the target. If subjects intended to minimize time also, then the effort-minimizing strategy would be a time-minimizing strategy from the subjects' point of view.

Nevertheless, the fact that subjects adopt the effort-minimizing strategy suggests that subjects' strategy may be characterized as the solution of a well-defined minimization problem, the same problem that the effort-minimizing IP solves.

The usefulness of the normalized performance measure strongly suggests that subjects' performance may be limited mainly by factors such as subjects' knowledge of the object's dynamics, which is the internal model that subjects develop and deploy. Alternative limitations on performance, such as errors in visual registration of object position or errors in timing, were ruled against by results from simulations of IPs whose visual or timing precision were intentionally degraded. To account for subjects' performance, errors in their vision or variability in their timing would have had to be so large as to be unrealistic.

The experimental task was designed to reduce a priori the parameters of motor input to a single one: button-pressing time. Similarly, the visual information was reduced to distance to target, vO 's position, velocity, and acceleration. This was done to enhance the meaningfulness of both IP analysis and the perceptual and motor noise models. Of course, there are many other noise models that could be used instead of, or in combination with, the ones explored here. In our selection of noise models, we attempted to use those who could be at least partially based on known experimental data. This helped us devise meaningful computational equivalents and their parameters, as well as a realistic range of variation of these parameters. Other potentially relevant sources of noise, such as variability in the perception of important components of vO 's movement (such as its smoothly varying velocity), particularly during action, have not, to the best of our knowledge, been investigated to date. This limits the usefulness of noise models of such potentially important processes. On the other hand, the approach described here could be used to examine precisely those processes, and their associated noise. For example, it could be possible to design a task in which most of the performance error could be attributed to noise in the perception of smooth velocity of objects during action, and use a specifically designed IP to assess performance.

An important possibility is that the internal model may contain abstract, function-like information about virtual object dynamics. Our data, though, do not rule out an alternative, namely that subjects used a tabular representation of input-output pairs [Atkeson 1989]. Such a table would associate, for example, (i) values of the object's starting distance from the target with (ii) corresponding button press durations. However, we propose that our methodology defines some basic tools that could be extended to explore this important and complex problem under novel experimental conditions.

We should stress that the internal model estimate produced by the dynamics-normalized measure does not assume that the internal model contains abstract information about the virtual object. We suggest that this estimate would be meaningful even if the internal model were totally or partially based on an input-output table similar to the ones discussed here. Further work is needed to evaluate how motor, perceptual, and cognitive factors influence this estimate, particularly early in training. For example, subjects' control strategy could change jointly with their internal model during training. Thus the dynamics-normalized measure would reflect this. We believe that this quantitative estimate

of the internal model is quite important in understanding the neural processes that support complex action-oriented behaviors. The linear models used to define the estimate (Appendix B), as well as the dynamics-normalized performance measure (Appendix A) can be generalized to include more powerful approaches, such as those of nonlinear system identification.

Where in the brain might an internal model of object dynamics be represented? The cerebellum [Miall 1998; Kawato and Wolpert 1998], dorsal premotor, and posterior parietal cortices [Shadmehr and Holcomb 1997] are among the sites in which internal models or components of such models are likeliest to be represented. Among these candidates, the cerebellum has attracted the greatest interest, in part because cerebellar subregions receive parallel inputs from regions involved in motor planning, and from regions that participate in visuomotor coordinate transformations [Kawato 1999]. However, nearly all empirical and theoretical treatments of internal models have focused on internal models of limbs, rather than of external objects, as in our study [Lackner and DiZio 1994; Conditt et al. 1997]. One recent functional neuroimaging study, though, did examine the acquisition of an internal model of an object. Imamizu et al. [2000] implicated an area near the cerebellum's posterior superior fissure in the acquisition of a new internal model of an external object, a rotationally remapped relationship between movements of a computer mouse and consequent movements of the cursor. Because Imamizu and colleagues presented subjects with a static, fixed transformation, an internal model that remapped motor behavior for that fixed transformation would not necessarily include objects' dynamics, or other time-dependent variables. As a result, although the cerebellum very well may be the depository of internal models for both limbs and objects, it cannot be said for certain that cerebellar circuits participate in generating an internal model of a time-varying (dynamical) external object. This applies to the virtual objects in our study, but it applies as well to real objects with which people interact in everyday life. Because it minimizes motor factors, the approach presented here is well suited for use with brain imaging techniques to help illuminate the neural substrates of the internal models of dynamical, external objects. In addition, the approach allows to correlate time-varying signals (such as fMRI's) to significant task events through the use of IP analysis.

The approach presented here affords a way to quantitatively estimate the characteristics of the internal model that a person uses while interacting with virtual objects. We propose that IPs could be used to model and quantify subjects' intended behavior during late (Appendix B) and early stages of visuomotor learning. This approach could also be extended to other experimental tasks, including tasks that involve complex, multimodal virtual objects. In addition to illuminating the characteristics of internal models, the performance measures and general approach introduced here could be extended to evaluate human operators' behavior in multimodal virtual environments. In particular, the specification of an IP could be used to identify the perceptual and motor factors that limit an operator's effectiveness.

APPENDICES

A. DESIGN OF THE DYNAMICS-NORMALIZED PERFORMANCE MEASURE

On any trial, an IP employs its perfect knowledge of the object's dynamics so as to move the object along a trajectory defined by

$$\text{Trajectory}_{\text{IP}} = h(t) * C_{\text{IP}} \quad (\text{A.1})$$

where $h(t)$ is the impulse response function of the virtual object, C_{IP} is the control signal generated by the IP in order to achieve optimal performance (minimizing either time or effort), and $*$ is the convolution operator [Spiegel 1974]. By definition, this trajectory will bring the vO into alignment with the target. Now, by replacing C_{IP} with the human subject's control signal for the same trial, C_{S} , in

Eq. (A.1) gives:

$$\text{Trajectory}_{\text{IP}} = h(t) * C_S \quad (\text{A.2})$$

But this equation holds true only if C_S is equal to C_{IP} . Substituting a function $h_N(t)$ for $h(t)$, Eq. (A.2) becomes

$$\text{Trajectory}_{\text{IP}} = h_N(t) * C_S \quad (\text{A.3})$$

Eq. (A.3) holds always, for a suitable function $h_N(t)$. Solving Eq. (A.3) for $h_N(t)$ gives

$$h_N = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}\{\text{Trajectory}_{\text{IP}}\}}{\mathcal{F}\{C_S\}} \right\} \quad (\text{A.4})$$

where \mathcal{F} and \mathcal{F}^{-1} are the Fourier and inverse Fourier transform operators, respectively [Spiegel 1974]. If C_S approaches C_{IP} , then $h_N(t) \rightarrow h(t)$, as Eq. (A.1) shows. This means that the closer $h_N(t)$ is to $h(t)$, the closer a subject's behavior will be to that of the IP. Note that this is true for *any* IP, not only for ones used in this study.

We can quantify the similarity of $h_N(t)$ and $h(t)$ in the following way. Substituting a Dirac delta function for $C_S(t)$ in Eq. (1), and integrating twice yields

$$m x(t) - \int_0^t (u(t') - b x(t')) dt' = 0 \quad (\text{A.5})$$

where $u(t)$ is the unit step function. Eq. (A.5) is satisfied if $x(t) = h(t)$ (virtual object's impulse response function, Eq. (2)) and leads to a dimensionless trialwise measure of error for $h_N(t)$:

$$\epsilon(h_N) = \frac{1}{t_{\text{dw}}} \sum_{j=1}^n | m h_N(t_j) - \int_0^{t_j} (u(t') - b h_N(t')) dt' | \quad (\text{A.6})$$

Eq. (A.6) evaluates the absolute value of Eq. (A.5) at discrete times t_j , and then cumulates those values over a data window of duration $t_{\text{dw}}=10$ s. As a subject's normalized performance $h_N(t)$ approaches $h(t)$, $\epsilon(h_N)$ decreases to a limit of zero, which is reached if $h(t) = h_N(t)$. We describe $\epsilon(h_N)$ as a "normalized measure" because with it a subject's performance is measured with respect to vO dynamics. Eqs. (A.4) and (A.6) were computed numerically using Matlab. The Fourier operators were approximated by using a fast Fourier transform with a 30-s data window [Press 1992]. For these calculations, $C_S(t)$ was subsampled at 250 Hz.

B. DYNAMICS-NORMALIZED PERFORMANCE AND INTERNAL MODELS OF VIRTUAL OBJECTS

By definition, $h(t)$ (the vO 's impulse response function, Eq. (2)) is the IP's internal model of the object's dynamics (Eq. (A.1), Appendix A). In the experiment described here, subjects approached the behavior of the effort-minimizing IP. Then, it is reasonable to use Eq. (A.1) as the basis for a description of the subject's behavior late in training. Following the form of Eq. (A.1), we can write:

$$\text{Trajectory}_S = h_i(t) * C_S \quad (\text{B.1})$$

where C_S is the subject's control signal intended to move the object along the desired path, Trajectory_S . The function $h_i(t)$, in analogy with function $h(t)$ in Eq. (A.1), would be interpreted as the subject's internal model of the virtual object. Our data confirm that human subjects perform much like an open-loop, predictive, effort-minimizing IP. This makes it possible to use the IP's $\text{Trajectory}_{\text{IP}}$ to model a human subject's intended trajectory, Trajectory_S , which produces

$$\text{Trajectory}_{\text{IP}} = h_i(t) * C_S \quad (\text{B.2})$$

which is identical to Eq. (A.3), the basis for our dynamics-normalized performance measure (Eq. (A.6), Appendix A).

ACKNOWLEDGMENTS

We thank Larry Abbott, Wilson Geisler, Michael Kahana, Paul DiZio, James Lackner, Eve Marder, and Luda Requadt for helpful comments on earlier versions of this manuscript. Research supported by CONACYT (National Council of Science and Technology, Mexico), the James S. McDonnell and Keck Foundations, the Fulbright Fellowship program, and AFOSR grant F49620-03-1-0376. Parts of this work were presented at a meeting of the Society for Neuroscience, and at the Boston University International Conference on Vision, Recognition, Action: Neural Models of Mind and Machine.

REFERENCES

- ALTES, R. A. 1989. An interpretation of cortical maps in echolocating bats. *J. Acoust. Soc. Am.* 85, 934–942.
- ATKESON, C. G. 1989. Learning arm kinematics and dynamics. *Ann. Rev. Neurosci.* 12, 157–183.
- BADDELEY, A., INGRAM, H. A., AND MIALL, R. C. 2003. System identification applied to a visuomotor task: Near-optimal human performance in a noisy, changing task. *J. Neurosci.* 23, 3066–3075.
- BARLOW, H. B. 1980. The absolute efficiency of perceptual decisions. *Philos. Trans. R. Soc. London, Series B, Biol. Sci.* 290, 71–82.
- BOGARTZ, R. S. 1971. Least squares methods for locating function shifts and disturbance regions. *Psychol. Bull.* 75, 294–296.
- BURBECK, C. A. 1986. Exposure-duration effects in localization judgments. *J. Opt. Soc. Am. A* 3, 1983–1988.
- BURBECK, C. A. 1987. Position and spatial frequency in large-scale localization judgments. *Vision Res.* 27, 417–427.
- BURBECK, C. A. AND YAP, Y. L. 1990. Two mechanisms for localization? Evidence for separation dependent and separation-independent processing of position information. *Vision Res.* 30, 739–750.
- CONDITT, M. A., GANDOLFO, F., AND MUSSA-IVALDI, F. 1997. The motor system does not learn the dynamics of the arm by rote memorization of past experience. *J. Neurophysiol.* 78, 554–560.
- DAVIDSON, P. R. AND WOLPERT, D. M. 2003. Motor learning and prediction in a variable environment. *Curr. Opin. Neurobiol.* 13, 1–6.
- DZHAFAROV, E. N., SEKULER, R., AND ALLIK, J. 1993. Detection of changes in speed and direction of motion: reaction time analysis. *Percept. Psychophys.* 54, 733–750.
- FLANAGAN, J. R., VETTER, P., JOHANSSON, R., AND WOLPERT, D. M. 2003. Prediction precedes control in motor learning. *Curr. Biol.* 13, 146–150.
- GEISLER, W. S. 1989. Sequential ideal-observer analysis of visual discriminations. *Psychol. Rev.* 96, 267–314.
- HOHNSBEIN, J. AND MATEEFF, S. 1998. The time it takes to detect changes in speed and direction of visual motion. *Vision Res.* 38, 17, 2569–2573.
- IMAMIZU, H., MIYAUCHI, S., TAMADA, T., SASAKI, Y., TAKINO, R., PUTZ, B., YOSHIOKA, T., AND KAWATO, M. 2000. Human cerebellar activity reflecting an acquired internal model of a new tool. *Nature* 403, 192–195.
- KAWATO, M. 1999. Internal models for motor control and trajectory planning. *Curr. Opin. Neurobiol.* 9, 718–727.
- KAWATO, M. AND WOLPERT, D. M. 1998. Internal models for motor control. *Sensory Guidance of Movement*. John Wiley, New York, 291–307.
- KEELE, S. W., POKORNY, R. A., CORCOS, D. M., AND IVRY, R. 1985. Do perception and motor production share common timing mechanisms? A correlational study. *Acta Psychol.* 60, 173–191.
- LACKNER, J. AND DiZIO, P. 1994. Rapid adaptation to Coriolis force perturbations of arm trajectory. *J. Neurophysiol.* 72, 299–313.
- McKEE, S. P. AND WATAMANIUK, S. N. J. 1994. The psychophysics of motion perception. *Visual Detection of Motion*. Academic Press, New York, 84–114.
- MIALL, R. C. 1998. The cerebellum, predictive control and motor coordination. *Sensory Guidance of Movement*. John Wiley, New York, 272–290.
- OGATA, K. 1978. *System Dynamics*. Prentice-Hall, Englewood Cliffs, NJ.
- PRESS, W. H. 1992. *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed. Cambridge University Press, New York.
- SHADMEHR, R. AND HOLCOMB, H. H. 1997. Neural correlates of motor memory consolidation. *Science* 277, 821–825.
- SPIEGEL, M. R. 1974. *Fourier Analysis*. McGraw-Hill Book Company, New York.

- WAUGH, S. J. AND LEVI, D. M. 1995. Spatial alignment across gaps: Contributions of orientation and spatial scale. *J. Opt. Soc. Am. A* 12, 2305–2317.
- WOLPERT, D. M. AND FLANAGAN, J. R. 2001. Motor prediction. *Curr. Biol.* 11, R729–R732.
- WOLPERT, D. M. AND GHARAMANI, Z. 2000. Computational principles of movement neuroscience. *Nature Neurosci. (Suppl.)* 38, 1212–1217.

Received July 2003; revised May 2004, July 2004; accepted July 2004