

Movement Planning With Probabilistic Target Information

Todd E. Hudson, Laurence T. Maloney, and Michael S. Landy

Department of Psychology and Center for Neural Science, New York University, New York, New York

Submitted 2 August 2007; accepted in final form 21 September 2007

Hudson TE, Maloney LT, Landy MS. Movement planning with probabilistic target information. *J Neurophysiol* 98: 3034–3046, 2007. First published September 26, 2007; doi:10.1152/jn.00858.2007. We examined how subjects plan speeded reaching movements when the precise target of the movement is not known at movement onset. Before each reach, subjects were given only a probability distribution on possible target positions. Only after completing part of the movement did the actual target appear. In separate experiments we varied the location of the mode and the scale of the prior distribution for possible targets. In both cases we found that subjects made use of prior probability information when planning reaches. We also devised two tests (Composite Benefit and Row Dominance tests) to determine whether subjects' performance met necessary conditions for optimality (defined as maximizing expected gain). We could not reject the hypothesis of optimality in the experiment where we varied the mode of the prior, but departures from optimality were found in response to changes in the scale of prior distributions.

INTRODUCTION

Performance in speeded-reaching tasks is often assessed by examining movements toward a spatial target at a known position in space. The target is visible before the start of movement and the key task for the motor system is to plan the most effective movement possible to reach the target (e.g., Körding and Wolpert 2004; Sabes and Jordan 1997; Todorov and Jordan 2002; Trommershäuser et al. 2003a,b). Other researchers have demonstrated that the motor system can update a planned movement in response to unanticipated changes in position, velocity, and visual properties of a fixed target (Brenner and Smeets 2004; Elliott et al. 1999; Komilis et al. 1993; Péllisson et al. 1986; Saunders and Knill 2004; Schmidt 2002). In all of these studies, a specific target is visible before movement onset even if the subject is fully aware it may unpredictably change location during the actual movement.

It is conceptually difficult to separate movement planning from movement execution in such tasks because the movement plan (including possible compensation for changes in target location) would likely be fully formed before movement onset (e.g., Bédard and Proteau 2004; Gribble et al. 2003; Heath et al. 2004; Rabin and Gordon 2004; Saunders and Knill 2004; Torres and Zipser 2004; Vindras and Viviani 2002). There are, however, natural movements for which there is substantial initial uncertainty concerning the final spatial goal of the movement, and the initial part of the movement must therefore be planned relative to the uncertainty of the goal information available before movement. In water polo, for example, an attacker must often plan and initiate a shot on the goal while a defender is simultaneously attempting to block the shot. Nei-

ther attacker nor defender can anticipate with certainty the actions of the other at movement onset and each can potentially react to the other's movement during the brief duration of the attack. The initial movement planning of either player should allow for a range of possible continuations that have a high probability of producing a successful outcome, each consistent with biophysical constraints imposed by the joints and the maximum torque-generating capabilities of the muscles. There will be an optimal initial trajectory that can be planned by the attacker based on (possibly) imperfect knowledge of the location of the goal, the biophysical limits of the motor system, prior information about the most likely defensive movements of the opponent, and the likelihood of hitting the goal given initial positions, velocities, accelerations, and so forth of the arm's initial trajectory.

Of course, whatever the attacker's eventual choice, the ultimate outcome of the chosen movement plan is to define the probability of success. The optimal movement plan would therefore be the one that maximizes this probability. In what follows, we have two major goals. The first is to test whether subjects are capable of modifying aim points, velocities, and so forth during the initial portion of a reach in response to probability information acquired before reach initiation. Given that this is the case, we will test subjects' performance to determine whether it meets necessary conditions for optimality (the Composite Benefit criterion, subsequently described, and the Row Dominance criterion, described in RESULTS), where optimal performance is defined to be performance that maximizes expected gain.

We will first describe the task and our theoretical framework for a simplified case. In this case, we present subjects with two possible targets (Fig. 1, gray rectangles). One of the targets is the correct target but, at the start of the movement, the subject does not know which. Once the subject's fingertip has traveled one third of the way to the target array (and passed through an invisible *trigger plane*, drawn as a dashed horizontal line in Fig. 1), the correct target is indicated visually and only after this point can the subject know with certainty which target carries a reward. The subject receives a reward by touching the correct target within 600 ms of movement onset and is penalized for slower (>600 ms) movements. This 600 ms includes the time needed to reach the trigger plane and also the time needed to travel from the trigger plane to the display screen containing the targets. Before the start of each trial, the available information defines the prior probabilities π_A and π_B that T_A or T_B is the correct target ($\pi_A + \pi_B = 1$). This prior probability distribution is all the target information that the subject has to plan the initial part of the movement from the

Address for reprint requests and other correspondence: T. E. Hudson, Department of Psychology, New York University, 6 Washington Place, 8th Floor, New York, NY 10003 (E-mail: hudson@cns.nyu.edu).

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.




FIG. 1. Thought experiment. Subject attempts to touch small targets on a screen. Movement must be completed within a short time limit (<600 ms). On each trial, there are 2 possible targets, T_A or T_B , one of which will be the actual target for that trial. At the start of the trial, the subject knows only that A will be the true target for that trial with probability π_A and B will be the actual target with probability π_B ($\pi_A + \pi_B = 1$). After completing part of movement to the screen, the subject learns which of the possible targets A or B is the actual target. If $\pi_A = 1$ then the subject knew that A would be the target on that trial and could simply plan a movement to A . Information as to target identity provided in mid-movement is redundant (similarly if $\pi_B = 1$). However, if, for example, π_A is 0.7 (and thus π_B is 0.3), then the subject may plan a “composite” movement plan that has 2 parts. Initial portion of the reach can be planned with certainty before movement initiation, here given by the trajectory indicated by the bold solid line. Because it depends on the target information gained when crossing the trigger plane, the final portion of the reach cannot be planned with certainty before reach initiation. Here, the final portion of the composite reach continues with either mean trajectory τ'_A or τ'_B after the identity of the target is known.

starting point to the trigger plane where the location of the target will be learned.

How should the ideal movement planner plan such a movement, particularly during the initial part of the movement up to the trigger plane? First, we consider special cases. Suppose that π_A is 1 (and therefore π_B is 0). Here, it is certain that T_A is the correct target and the subject can simply plan an optimal movement to T_A (a *determinate target*). We refer to the outcome of movement planning as a movement plan or strategy, denoted s . The ideal movement planner should adopt a movement plan s_A that leads to a mean spatial trajectory such as the one labeled τ_A that ends at T_A and that maximizes the probability of reaching the target within 600 ms and earning the reward. We refer to such a plan as a *simple movement plan* and the resulting reaches and trajectories as *simple reaches* and *simple movement trajectories*. A movement planner must specify not just a spatial trajectory but also how the trajectory evolves across time. For simplicity in presentation, however, we defer discussion of planning movement velocity or higher temporal derivatives.

We denote the probability of acquiring T_A (a hit on T_A , denoted H_A) with this simple movement plan s_A as $p(H_A | s_A)$. This is the probability of earning the reward with this trajectory. The ideal movement planner would pick the simple

movement plan that maximizes this probability. Similarly, if the ideal movement planner knew that the correct target was T_B at the beginning of the trial, then a simple movement plan s_B would be adopted, leading to a mean trajectory τ_B terminating at T_B . The probability of earning the reward with this plan is $p(H_B | s_B)$. The two mean trajectories corresponding to these two simple movement plans are marked by dashed curves in Fig. 1.

In either of these cases, the subject is simply asked to optimize movement to a determinate target. In particular, the target information that the subject receives after passing through the trigger plane is redundant and the optimal movement planner will ignore the trigger plane and the alternative target location in planning simple movements to determinate targets.

Suppose now that the ideal movement planner is told that π_A is 0.7 (and therefore $\pi_B = 0.3$); these are *probabilistic targets*. When the hand passes through the trigger plane, either T_A or T_B will be revealed as the actual target. How should the ideal movement planner plan the resulting composite movement s (where a composite movement is one that has multiple possible completions, in this case one that can be completed toward T_A or T_B , with mean trajectory τ'_A or τ'_B ; Fig. 1)?¹ In particular, how should the subject plan the initial phase of the composite movement (extending up to the trigger plane)? A movement planner could simply plan a simple movement s_A to T_A (with mean trajectory τ_A), given that T_A is more likely, ignoring the information available at the trigger plane and ignoring T_B even when it is the correct target. The probability of earning the reward (acquiring target T_A) is then $p(H_A | s_A)\pi_A$. Alternatively, the initial trajectory to the trigger plane could be planned such that it intersects the trigger plane between the intersection points of the simple trajectories to T_A and T_B . The solid trajectory in Fig. 1 illustrates a possible initial trajectory intersecting the trigger plane at such an intermediate location. The initial portion of the composite movement can continue as trajectory τ'_A to T_A or trajectory τ'_B to T_B , both drawn as solid lines in Fig. 1. As drawn, the trajectory of the composite movement leading to T_A deviates less from the optimal simple trajectory to T_A , reflecting the possibility that the subject may choose to favor the more likely target. In planning this trajectory, the movement planner has to allow for the cost (if any) of registering the correct target information at the trigger plane and the cost of updating the movement plan to now move toward the correct target.

Composite benefit criterion

The ultimate consequence of choosing a composite movement plan s is to affect the probabilities of hitting either target, T_A or T_B , when it is the correct continuation of the initial portion of s . We denote by $p(H_A | s \cdot T_A)$ the probability of hitting target T_A on trials in which T_A is the target and composite movement plan s is used; we similarly define $p(H_B | s \cdot T_B)$. We also assume that there is zero probability of hitting nontarget T_A when T_B is the target [i.e., $p(H_A | s \cdot T_B) = 0$], and vice versa, given that when the true target is revealed

¹ To be clear, a composite movement plan is a composite in the sense that it is composed of an initial phase and an end phase, where the end phase cannot be planned with certainty until the initial phase is completed. It is *not* a weighted mixture or superposition of simple movement plans.

the others are removed from the display. Then, the overall probability of earning a reward on each trial is

$$p(\text{Reward}|s) = p(H_A|s \cdot T_A)\pi_A + p(H_B|s \cdot T_B)\pi_B \quad (1)$$

Note that $p(H_A|s \cdot T_A)$ and $p(H_B|s \cdot T_B)$ need not sum to 1; a subject could perform poorly both when T_A is the target and also when T_B is the target. We expect that $p(H_A|s \cdot T_A) \geq p(H_B|s \cdot T_B)$ in our example because the former is the probability of hitting T_A with a movement plan likely to be biased toward the more probable target location, T_A . A composite plan s is preferable to executing a simple plan to one of the two targets when

$$p(H_A|s \cdot T_A)\pi_A + p(H_B|s \cdot T_B)\pi_B > \max \{ p(H_A|s_A \cdot T_A)\pi_A, p(H_B|s_B \cdot T_B)\pi_B \} \quad (2)$$

In the experiments we report, we will use N targets. Equation 1 then becomes

$$p(\text{Reward}|s) = \sum_{j=1}^N p(H_j|s \cdot T_j)\pi_j \quad (3)$$

and the composite plan s is preferable to any simple plan only when

$$\sum_{j=1}^N p(H_j|s \cdot T_j)\pi_j > \max \{ p(H_1|s_1 \cdot T_1)\pi_1, \dots, p(H_N|s_N \cdot T_N)\pi_N \} \quad (4)$$

We refer to the condition defined by Eq. 4 as the Composite Benefit criterion. It is a necessary condition for optimal movement planning.

Composite-movement planning

The spatial trajectory is not the only aspect of the movement plan that the planner can consider in formulating a composite movement plan. Recall that the reach must be completed within 600 ms of its initiation or no reward is earned and a penalty is imposed. Given that, with the composite plan s , the subject must either accelerate left to T_A or right to T_B after reaching the trigger plane, it may be preferable to reach the trigger plane traveling at a lower speed than if following either of the simple trajectories to reduce the torques required to accomplish trajectory adjustments. However, taking a longer time to reach the trigger plane or passing through the trigger plane at low speed eats into the time available to complete the movement, and the optimal trade-off between reduced time and the noise from increased torque production is likely to be complex. Regardless of the details of the trade-off, subjects may choose to vary not only the spatial path but also the velocity profile of the path to obtain the highest total reward possible.

What does the motor system plan when planning composite movements? The initial phase of the reach cannot be programmed as a function of the location of a target as is typically assumed (e.g., Abrams et al. 1990; Woodworth 1899) but the subject can plan the *motor state* of the fingertip (location, orientation, velocity, acceleration, etc.) as it passes through the trigger plane. We seek to determine whether and how the subject alters this planned motor state in response to changes in the prior probabilities of the targets. We discuss next the possible responses to specific changes in the prior distribution.

Predictions

In planning the initial part of a movement, we expect subjects to select both a movement goal for the initial movement and a suitable control law for its implementation. In our task, subjects cannot plan an optimal reach to the unknown target location before movement onset. However, they can plan the initial portion of the reach to produce a state of the motor system at the trigger plane that is maximally advantageous for later acquisition of the target (and reward) once it is known. We do not know the form of the initial movement plan and the interpretation of our experiments does not require this knowledge. Participants may initially plan a movement only through the trigger plane or they may choose an initial goal location on the display screen and an intended speed of movement, and change that goal after the target is displayed. Whatever the form and goal of the initial plan, the plan and its implementation determine the state of the fingertip when it passes through the trigger plane (its position, velocity, acceleration, etc.), and it is these kinematic variables that we measure and relate to performance in the task. The state of the fingertip at the trigger plane determines the probability of subsequent target acquisition (thereby also determining expected gain). We will therefore equate the outcome of movement planning with consistent, patterned changes in the state of the fingertip at the trigger plane.

We make two conjectures concerning how an ideal subject will perform.

1 Changes in the location of the highest probability target should serve to shift the location at which the fingertip passes through the trigger plane. If the fingertip passes through the trigger plane at a horizontal location c_x when the highest probability target is at the center of the set of possible target locations, then a leftward/rightward shift in the location of the highest probability target would shift c_x leftward/rightward. This possibility is tested by providing subjects in a first (“*Location*”) experiment with a series of probability distributions that differ in the location of their mode. Because of the complexity of our task, we cannot compute the optimal movement plan. Yet, it is possible to test whether human performance is consistent with an optimal solution using a test based on the Composite Benefit criterion described earlier, and a second test based on an additional necessary condition for optimality, the Row Dominance criterion, subsequently explained. We refer to the tests based on these two criteria as the Composite Benefit and the Row Dominance tests, respectively. Analogues of these tests, particularly the latter, should be useful in comparing human to ideal performance in a wide variety of movement tasks where generating precise predictions of quantitatively optimal performance is infeasible, given the complexities of modeling movement trajectories under biomechanical constraints and neural limitations of the motor system that are not fully understood.

2 Reaching the same point on the trigger plane but at reduced speed, for example, might be a proper response to an increase of uncertainty in the location of the target because high velocities at the trigger plane mean that any trajectory change will result in increased torques and increased movement error in generating the motor commands needed to change direction (Hamilton et al. 2004; Todorov 2002). We

will investigate this possibility in a second (“*Scale*”) experiment where we vary the width of the prior distribution, leaving the location of the mode unchanged.

To anticipate, we found that subjects modified position and velocity in the two experiments, respectively, in a manner consistent with the preceding qualitative predictions. In the Location experiment, where we varied the mode of the prior distribution, we could not reject the hypothesis of optimal movement planning by either of the two criteria considered. However, we did demonstrate that performance was suboptimal in the Scale experiment; subjects failed both the Composite Benefit test and the Row Dominance test. We discuss these results in relation to previous work demonstrating predictive control as well as recent work on Bayesian optimality in motor planning.

METHODS

Apparatus

Subjects sat at a custom-made (Mica-Tron) aluminum table that securely held a computer monitor behind a 43×61 -cm sheet of transparent polycarbonate. Stimuli were presented on a Sony Multi-Scan G500 with a functional display area of approximately 39.2×28.75 cm and pixels separated by 0.2 mm.

A Northern Digital Optotrak 3D motion capture system (with two three-camera heads) was used to measure the position of the subject’s right index finger, target screen, and tabletop with a set of eight infrared light-emitting-diode (IRED) markers (sampling rate: 200 Hz with IREDs strobed at 2,500 Hz). Four of the markers were embedded in the transparent polycarbonate screen that covered the computer monitor and allowed localization of the screen and integration of Optotrak and computer monitor frames of reference. The monitor reference frame was identified with the frontal x - z plane of the subject. A fifth marker was placed at the near edge of the tabletop to mark the start position of the reaches. The remaining three markers were attached to an extended ring that fitted over the distal joint of the subject’s right index finger. Optotrak measurements for these three markers were used to compute the location of a “virtual marker” at the tip of the finger (see *Protocol*). We calibrated the Optotrak cameras spatially before each experimental run, providing root-mean-square accuracy of 0.1 mm within the volume immediately surrounding the subject and monitor apparatus (~ 2 m³). Four IRED markers were embedded at precisely measured locations in the polycarbonate sheet to aid in registering the monitor within the Optotrak system before each experimental session. An additional IRED located at the front edge of the table marked the starting point for subjects’ movements. The experiment was run using Psychophysics Toolbox software (Brainard 1997; Pelli 1997) and the Northern Digital software library (for controlling the Optotrak) on a Pentium III Dell Precision workstation.

Targets

Possible target locations were represented as vertical bars on the screen. Bars were 32 pixels wide and 200 pixels high, about $24 \text{ min} \times 5 \text{ deg}$ at the subject’s viewing distance of 42.5 cm. Each bar was partitioned into 100 (4×25) segments, colored either light or dark gray, and presented against a black background. The relative number of light segments indicated the probability of that bar’s containing the target (Fig. 2).

Prior probability distributions

In the Location experiment, we used five prior probability distributions defined on nine equispaced targets (Fig. 2A). One of the five

central bars (the third bar in Fig. 2A) had prior probability 0.68 of being the target, whereas the remaining bars each had probability 0.04. In effect, we varied the location of the mode of the probability distribution while keeping its width constant.

In the Scale experiment, we used three prior probability distributions defined on seven equispaced targets (Fig. 2B). Each probability distribution was spatially symmetric, with its maximum extended over one to five of seven possible target locations, and the remaining probability mass distributed evenly in the tails of the distribution. These probability mass functions will be referred to as the 1) low-, 2) medium-, and 3) high-certainty conditions, for which the prior probabilities of each of the target locations are as follows:

- 1 Low-certainty: $\pi = [0.075 \ 0.17 \ 0.17 \ 0.17 \ 0.17 \ 0.17 \ 0.075]$
- 2 Medium-certainty: $\pi = [0.025 \ 0.025 \ 0.3 \ 0.3 \ 0.3 \ 0.025 \ 0.025]$
- 3 High-certainty: $\pi = [0.025 \ 0.025 \ 0.025 \ 0.85 \ 0.025 \ 0.025 \ 0.025]$

We can quantify the uncertainty associated with each prior by its Shannon entropy in bits, calculated as $H(\pi) = -\sum_i \pi_i \log_2 \pi_i$. These values were 2.73, 2.10, and 1.00 for the high-, medium-, and low-certainty conditions, respectively. In contrast, the entropy was 1.86 for all priors in the Location experiment.

Protocol

A key comparison to be made in these studies is between reaches to identical targets made under certain and uncertain information—that is, between simple and composite reaches to the same target locations. For this reason, each subject’s experimental session began with a series of reaches made toward the same target locations and prior distributions as described earlier, but with the correct target location indicated before each reach. These determinate targets occurred at the various locations within each distribution with the same frequency as indicated by the probability distribution. Target locations during simple reaches were indicated before reach initiation by a pair of small gray dots flanking the correct potential target bar; at the trigger plane, the nontarget bars disappeared, leaving just the target (now colored entirely white). Subjects were aware of the visual coding of prior probabilities by small white squares within the target bars, and these reaches gave subjects a separate opportunity to learn the frequencies with which each bar’s location would become the target for each of the probability distributions, while they simultaneously made simple reaches to known target locations. After initial reaches to determinate targets, subjects were instructed that they would be pointing to the same targets on the screen without the indicator dots and paid a bonus based on the sum of the point values they earned in each trial during this “test” phase of the experiment. The subject earned 15 points for hitting the target, lost no points for missing the target, and lost 45 points for reaching the screen after the time-out period.

The Location and Scale experiments consisted of a single session of 250 or 300 simple reaches, followed by 625 or 600 composite reaches, respectively. Before the experiment, there was a calibration sequence in which the three markers held on the ring were calibrated to the position of the fingertip. The calibration procedure consisted of placing the tip of the right index finger over the center of one of the IREDs embedded in the screen while recording the locations of the three ring markers, to compare with the known location of the screen marker. The location of the virtual fingertip position could then be calculated on-line during the experiment from the calibration information and the current positions of the three ring markers. Subjects were told they could rest at any time between reaches to avoid fatigue. Subjects never waited more than a few seconds between reaches.

Sequence of events within a trial

The following are characteristic of reaches to all targets, determinate and probabilistic: At the beginning of each reach, the




FIG. 2. Stimuli. *A*: stimuli in the Location experiment consisted of 9 bars. Proportion of square white elements in each bar was equal to the probability that bar would eventually become the movement target. High-probability bar could be at any of the central 5 positions. An example of the high probability located at the third bar is shown. *B*: in the Scale experiment, 7 bars were used. Either 1, 3, or 5 bars had probability higher than that of the others. An example of the medium-certainty condition (3 higher-probability bars) is shown.

finger tip was positioned at a start location, 350 mm in front of the screen and 1.5 mm to the right of the screen center. This position was indicated to the subject as the intersection of the edge of the custom tabletop and a raised ridge orthogonal to the tabletop edge. When the fingertip crossed a virtual frontal plane 348 mm in front of the screen while returning to the start position, the prior distribution for the next trial was signaled by an auditory cue and, after a 1-s pause, the visual representation of the prior was presented on the screen. This probability distribution for possible target locations was positioned near the center of the screen, jittered to the left or right by a maximum of ± 1.6 cm (randomly drawn from a uniform distribution). At any time after the presentation of the prior on the screen, the subject could begin the reach. The timer began as the fingertip recrossed the virtual frontal plane 348 mm in front of the screen.

When the fingertip crossed a second virtual plane (the trigger plane) located one third of the distance to the screen (232 mm in front of the screen), the target was triggered and the visual representation of the prior probability density was replaced by a single white bar at the true target location.

The reach terminated when the fingertip crossed a third virtual frontal plane 3 mm in front of the screen and the fingertip velocity fell to < 1 mm/s. Three distinct auditory indicators were used to signal whether the subject had hit or missed the target, or whether the movement had been too slow. In addition, the words "HIT," "MISS," or "TOO SLOW" were displayed after termination of the movement. Feedback was displayed until the fingertip returned to the start position, behind the first virtual plane 348 mm in front of the screen. Returning to the start position began the next trial and the screen was momentarily blanked.

Differences between reaches to determinate and probabilistic targets

The main difference between reaches to determinate and probabilistic target locations was that the true target locations were displayed before reach initiation for simple reaches to determinate targets, but not for composite reaches to probabilistic targets. This was accomplished by displaying two small, low-contrast circles on either side of the center of the bar that was to become the target (before the reach). This provided subjects with perfect information about target location while simultaneously allowing them to experience the frequency with which each bar became the target for each prior probability distribution. It was this experience with the frequency at which each location became the target that allowed subjects to learn each of the prior probability distributions.

Additional information concerning the timing of the reach and fingertip placement at the screen was also available during simple reaches. The proportion of total time elapsed during each reach to determinate targets was displayed as a timer bar, which provided an on-line indication of the time elapsed during the reach. The movement endpoint was displayed after each simple reach as a long thin vertical line whose vertical extent was greater than that of the target bars. This fingertip endpoint indicator was colored green for hits and red for misses. No fingertip endpoint indicator was presented when subjects timed out. Both the timer bar and the fingertip endpoint indicator were displayed until the screen was blanked and a new trial begun. The scatter of fingertip endpoints around the center of the target measured during simple reaches was used to determine the width of bar that would have produced 65% (Location experiment) or 85% (Scale experiment) hits. Although the visual representation of the bars

remained constant for all reaches, fingertip endpoints during reaches to probabilistic targets were rewarded only when they fell within the above-calculated distance from the center of the target bar. Naïve subjects did not detect this manipulation, which helped normalize performance across subjects. During execution of composite reaches, the bonus associated with the outcome of the current reach (15, 0, -45, for a hit, miss, or time-out, respectively) and a running total bonus score were displayed after each reach.

The first few reaches to determinate targets were typically less accurate because subjects were unfamiliar with the experimental apparatus and with making timed reaches. Only trials collected after performance had stabilized were used in later data analyses. We estimated performance across time as the probability of hitting the target in the immediately preceding 30 trials. We estimated asymptotic performance as the mean and SD of 30-point performance measures for the second half of the determinate trials. We discarded initial determinate trials until performance was within 2.5 SD of final performance. This resulted in removal of 46 of 1,500 determinate trials in the Location experiment and 89 of 1,800 in the Scale experiment. No conclusions are changed by inclusion/exclusion of these trials.

By measuring reaches to both determinate and probabilistic targets using the same prior probability distributions, we will be able to compare simple and composite reach trajectories to the same set of target locations, under the corresponding difference in uncertainty inherent in reaches to determinate and probabilistic target locations.

Subjects

In the Location experiment, subjects were between 19 and 34 yr of age, three male and three female; in the Scale experiment, subjects were between 23 and 34 yr of age, four male and two female. All subjects used the right hand for reaches in the experiment, although one (SG, Location experiment) uses her left hand for some tasks, including writing.

Data analysis

Several of the results presented here involve model comparison of nonnested models and are best analyzed with Bayesian methods (see Supplement²). Analyses are presented in detail in RESULTS to facilitate understanding of the rationale and advantages of each technique to the specific inference to be drawn from the data. Where appropriate, we present the results of standard statistical and likelihood-based methods for comparison.

Unlike a standard analysis, a Bayesian analysis requires not only a likelihood function, but also a prior probability distribution. We use Jeffreys priors in all Bayesian analyses. A Jeffreys prior corresponds to the weakest possible assumptions that we can make about model parameters and is commonly used in such analyses (Jaynes 2003; Jeffreys 1946).

RESULTS

Location experiment

In what follows, the *z*-dimension (height) is of little importance because the targets were elongated vertically and only the *x*-component of the fingertip position at the screen affected the outcome of a trial. We first projected the reach trajectories onto the tabletop and then calculated space-averaged trajectories along the *y*-axis (i.e., the average *x*-position as a function of *y*) for the central five target locations (determinate targets) and corresponding five conditions (probabilistic targets). The *x*-position of the fingertip at the trigger plane for reaches made to

the central five determinate target locations was compared with the *x*-position of the fingertip for reaches made to probabilistic targets in the five conditions with the corresponding peak probability locations.

Figure 3A shows space-averaged trajectories for each of the central five determinate targets (mean of 50 reaches/subject), as well as initial trajectories for the five probabilistic target conditions (mean of 125 reaches/subject). Although all trajectories are used in our analyses, the composite reaches shown in Fig. 3A continue from the trigger plane with averages only over reaches to the high-probability target location (mean of 85 reaches/subject) to reduce the complexity of the figure.

In addition to calculating the space-averaged trajectories shown in Fig. 3A, we determined whether there were significant carryover effects from one reach to the next on subjects' trigger plane crossing points. In other words, we asked whether a crossing point slightly to one side of average for a given reach would be followed by a correction to the same or the opposite side on one or more of the immediately subsequent

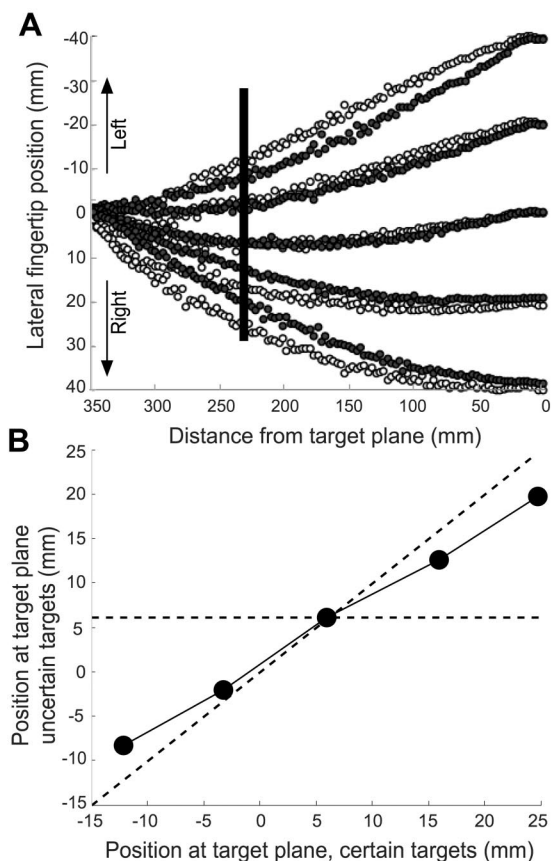


FIG. 3. Mean spatial trajectories and trigger plane positions in the Location experiment. *A*: mean lateral position as a function of distance from the screen (open symbols: determinate targets; filled symbols: reaches to the most probable of the probabilistic targets). Black vertical bar indicates the trigger plane. *B*: trigger plane crossing points for uncertain vs. certain target locations. Lateral distance from the center of the central target bar where the finger crossed the trigger plane during reaches to probabilistic targets plotted as a function of the same position during reaches to determinate targets; the 45° dashed line indicates the identity function, i.e., the expected outcome if the subjects adopted the same movement plan for uncertain targets as they had for certain targets. Horizontal dashed line indicates the expected outcome if subjects ignored the information contained in the probability distribution. In both plots, means are over conditions and subjects.

² The online version of this article contains supplemental data.

reaches. There were no significant autocorrelations of trigger plane crossing points beyond lag 0, indicating that the position at which subjects' fingertips crossed the trigger plane on a given trial was unaffected by the crossing points experienced in previous trials. This is perhaps an unsurprising result because the prior distributions were presented in an interleaved, unpredictable order.

Reach trajectories exhibited the characteristic slight curvature reported in other studies (e.g., Flanagan and Rao 1995; Goodbody and Wolpert 1999; Osu et al. 1997). The slight curvature seen in trajectories that did not require a large mid-reach adjustment (Fig. 3A) corresponded to a roughly constant rate of change of angular direction over the main body of the reach (discussed in the following text; see also Fig. 4 for similar results from the Scale experiment).

The increased uncertainty of reaches to probabilistic targets relative to determinate targets influenced the initial composite reach trajectories and was expected to produce a compression of the former's lateral trigger-plane crossing points relative to the simple-trajectory crossing points measured during reaches to determinate targets. However, because there was still substantial information concerning target location in each of the prior probability distributions, we expected the crossing points

of composite reaches to be biased in the direction of the location of the peak probability location, and therefore predict a slope between 0 and 1 when trigger-plane crossing points from simple trajectories are plotted against those from composite reach trajectories. Consequently, we were interested in determining whether a slope of $a = 1$ (no compression of crossing points), $0 < a < 1$ (partial compression), or $a = 0$ (full compression) captured the relationship between fingertip position at the trigger plane for simple trajectories to the central five targets and composite trajectories in the five test conditions.

Figure 3B shows the relationship between these crossing points for simple and composite reaches. The regression of average composite-reach trigger-plane crossing in the five test conditions on simple-reach trigger-plane crossing points for reach trajectories to the central five target locations had a least-squares fitted slope of $a = 0.760$. We can reject the hypotheses that the slope is 0 ($t = 77.1$; $P < 0.001$) or, separately, that it is 1 ($t = -24.3$; $P < 0.001$).

Although the preceding t -tests are the standard statistical tests for determining whether a slope is not 0 or 1, they do not provide a simultaneous test of the three hypotheses (no compression, partial compression, full compression) that takes into account the fact that there are many more possible slope values that are consistent with partial compression than with the other two alternatives. A better test of these hypotheses is possible when the probabilities of models incorporating the constraints that the slope is 0, 1, and between 0 and 1, respectively, are compared directly to one another. These probabilities automatically encode the discrepant numbers of possible slope values that are consistent with the three competing hypotheses. The probabilities of the three models were converted into odds ratios and these ratios were converted into a decibel measure, called *evidence*³ (Jaynes 2003). The evidence in decibels for full compression relative to the other two hypotheses is -82.1 dB. The evidence for zero compression is -17.3 dB and the evidence for partial compression is 23.3 dB. There is clearly more evidence for the hypothesis that the slope is strictly between 0 and 1 than for slope values of precisely zero or one.⁴

If either full or zero compression had been the preferred model, we would expect that the corresponding slope of 0 or 1 would be the best (highest-probability) estimate of the slope. However, given that partial compression was the preferred model ($y = ax$; $0 < a < 1$), we next calculated the posterior probability distribution associated with the range of possible slopes consistent with partial compression and the data, using an uninformative Jeffreys prior (Jeffreys 1946) for slopes. This distribution has its maximum at $a = 0.785$, close to the least-squares estimate of 0.760 reported earlier.

Haruno and colleagues (2001) described a model of motor control, MOSAIC, that provides for multiple controllers. At any instant, each controller suggests a motor command; these commands are weighted based on a set of "responsibility predictors." One can imagine an application of this model to the current experiment wherein one controller is associated with each potential target and, when invoked alone, produces

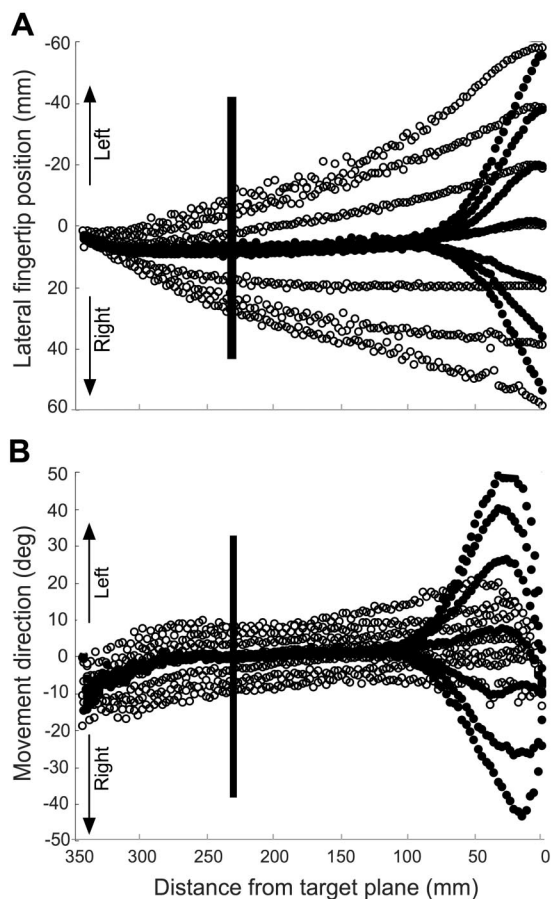


FIG. 4. Mean spatial and directional trajectories in the Scale experiment. *A*: mean lateral position as a function of distance from the screen (open symbols: determinate targets; filled symbols: probabilistic targets). *B*: mean movement angle (projected onto the horizontal plane) as a function of distance from the screen (open symbols: determinate targets; filled symbols: probabilistic targets). In both plots, means are over conditions and subjects. Black vertical bars indicate the trigger plane.

³ Positive evidence provides support for the hypothesis being tested and negative evidence provides support for the negation of that hypothesis, relative to the other hypothesis or set of hypotheses being tested.

⁴ Although as arbitrary as any significance threshold using P -values, we use a criterion for evidence of 3 dB, which corresponds to odds of nearly 2:1.

the simple trajectory to that target. In MOSAIC, these responsibility coefficients are learned based on forward-model prediction errors. However, consider a modification of MOSAIC for our probabilistic-target conditions in which the responsibility coefficients are equal to the corresponding target probabilities. This modified model predicts a mean composite trajectory equal to the target-probability-weighted average of the simple trajectories; that would predict partial compression with a slope of 0.64 and an intercept of 2 mm. The evidence favors the hypothesis of partial compression over this “mixtures-of-strategies” hypothesis by 3.7 dB. The mixtures-of-strategies hypothesis is also rejected by *t*-tests comparing the slope (0.64) and intercept (2 mm) predicted by a mixture of strategies to the best-fit slope (0.76, $P < 0.01$) and intercept (0.84 mm, $P < 0.01$). Thus we must reject this “mixtures-of-strategies” model for our Location experiment data.

ROW DOMINANCE TEST. We next tested whether subjects traded off accuracy at hitting low-probability targets for improved accuracy at hitting the same targets when they have high probability. We can test relative effectiveness of the observed initial reach trajectories by comparing the points earned by the subject in, say, condition 1 (leftmost high-probability target, with target prior probability distribution π_1 using the observed strategy s_1) with the expected number of points the subject would have earned had the subject instead used the strategy displayed in another condition (e.g., strategy s_2 from condition 2).

Each of the $k = [1, 2, \dots, 5]$ conditions in the experiment corresponded to a prior on the nine targets that we denote by the row vector $\pi_k = [\pi_{1k}, \dots, \pi_{9k}]$. Let $\mathbf{p}_k = [p_{1k}, \dots, p_{9k}]$ denote the frequency at which subjects hit each of the nine targets when each was the target while using the movement strategy adopted for condition k ; that is, $p_{ik} = p(H_i | s_k \cdot T_i)$. For example, the initial trajectory observed in the condition with the mode at the center target position resulted in hit frequencies at each of the nine targets of $\hat{\mathbf{p}}_3 = [0.267, 0.300, 0.233, 0.667, 0.708, 0.567, 0.167, 0.233, 0.133]$ based on the data. Clearly, this initial trajectory is much more effective in acquiring the central (5th) target position than, say, the 7th position.

The inner product $\langle \mathbf{p}_k, \pi_k \rangle = \sum_i p_{ik} \pi_{ik}$ is the sum of the prior for each target multiplied by the frequency at which that target was hit. That is, it is the expected hit rate when adopting strategy s_k in condition k . This expected hit rate is also proportional to the subject’s expected earnings in condition k using strategy s_k .

But what if the subject had used the movement strategy used in a different condition k' in condition k ? The subject’s rate of success would then be $\langle \mathbf{p}_{k'}, \pi_k \rangle$. If $\langle \mathbf{p}_{k'}, \pi_k \rangle \leq \langle \mathbf{p}_k, \pi_k \rangle$, then this alternative strategy for condition k would have earned less on average than the actual strategy used. That outcome is consistent with the claim that the subject has chosen the optimal movement strategy that this subject is capable of in condition k . However, if $\langle \mathbf{p}_{k'}, \pi_k \rangle > \langle \mathbf{p}_k, \pi_k \rangle$, we can reject this claim of optimality: the subject is capable of a movement strategy, exhibited in condition k' , that would have earned more in condition k than the strategy actually used.

We can compute the inner products of all pairings of hit-probability vectors $\mathbf{p}_{k'}$, $k' = 1, \dots, 5$ and priors π_k , $k = 1, \dots, 5$ as a 5×5 matrix and examine the match between movement strategy and prior. These are shown in Table 1. The k th row records the performance of each of the movement strategies k' in condition k (with prior π_k). The third row, for example, records how each of the movement strategies would have fared with prior π_3 . The maximum value is 0.584 (paired with \mathbf{p}_3 for strategy 3) and the minimum value is 0.401 (paired with \mathbf{p}_5 for strategy 5). Among the strategies evoked across conditions, the strategy chosen in condition 3 maximizes expected earnings in condition 3. A necessary condition for optimal performance (maximizing expected gain) is that the diagonal value in each row not be significantly less than any of the other entries in the row. This condition must hold for each row and we therefore call it the *Row Dominance criterion* and the corresponding test the *Row Dominance test*.

In the results summarized in Table 1, the diagonal entry in each row is greater than the other entries in the same row, not less, and therefore not significantly less (all P values for comparisons of row entries are >0.5). We do not reject the hypothesis of Row Dominance.

One objection to this test concerns its power. Suppose that, across the range of experimental conditions, subjects’ winnings are scarcely affected by picking the wrong movement plan and the outcome of the Row Dominance test simply captures this insensitivity. We can test a stronger claim than Row Dominance, that each diagonal entry is not only greater than or equal to the other entries in its row, but that the inequality is strict. That is, not only did the subjects pick a movement strategy that did not perform worse than another observed strategy, but had they used any of these movement strategies used for the other priors, they would have done significantly less well on average.

TABLE 1. Location Experiment: Row Dominance (pooled over subjects)

	Strategy 1	Strategy 2	Strategy 3	Strategy 4	Strategy 5
Condition 1	0.51 (0)	0.47 (14.0)	0.28 (203.8)	0.35 (104.2)	0.38 (74.2)
Condition 2	0.47 (45.9)	0.58 (0)	0.56 (5.8)	0.48 (46.1)	0.29 (287.3)
Condition 3	0.41 (114.4)	0.56 (7.5)	0.58 (0)	0.46 (67.5)	0.40 (125.6)
Condition 4	0.28 (279.5)	0.49 (28.3)	0.49 (28.2)	0.56 (0)	0.47 (40.1)
Condition 5	0.15 (442.7)	0.37 (61.6)	0.24 (235.2)	0.48 (2.4)	0.49 (0)

Results of the Row Dominance test for the Location experiment. We test whether subjects could have earned more on average in each experimental condition by using a strategy used in a different condition. The Bayesian optimal movement planner is one who selects the movement strategy for each condition that maximizes expected gain. If we find that our subjects could have done better in a condition by adopting the strategy they used in another condition, then we can reject the hypothesis of optimality. The entry in each cell of the table is the probability of obtaining a hit when one of the five observed reach strategies (indexed by column) is combined with one of the five prior distributions (indexed by row); values in parentheses report an evidence measure (dB) testing whether the corresponding probability is smaller than the probability on the main diagonal in the same row. Positive values are evidence that it is; negative values, that it is not (italicized entries are above the 3 dB threshold). Boldface entries mark the maximum of each row. Probabilities were calculated from the lumped data, with each subject contributing an equal number of trials.

We therefore tested this Strict Row Dominance hypothesis by calculating the probability that the values along the main diagonal were strictly greater than other values in the same row. The evidence values associated with this hypothesis (in dB) calculated from these probabilities are given in parentheses to the right of the expected hit rates (Table 1, calculated from hit frequencies pooled over all subjects; also see Supplement for confidence intervals surrounding estimates of expected hit rates). Positive evidence values⁵ favor the hypothesis that diagonal elements are strictly greater than the relevant off-diagonal element within that row, consistent with our prediction.

All maximum expected hit rates for each row occur along the main diagonal, consistent with our prediction. Italicized expected hit rates are below the diagonal elements by ≥ 3 dB. In this experiment, all off-diagonal rates are significantly below those on the diagonal except for the last comparison in the fifth row, which is just below the 3 dB criterion.

COMPOSITE BENEFIT TEST. In addition to testing Row Dominance, we can assess whether reach planning was consistent with a second necessary condition for optimality, the Composite Benefit criterion (Eq. 4). Equation 4 implies that an optimal reach planner will choose a simple movement plan to a single target, ignoring other possible targets and the information provided when crossing the trigger plane, when the expected hit rate using a simple movement plan for that target is greater than the overall expected hit rate for the composite movement plan. If a simple movement plan had been used to generate reaches in the Location experiment, a maximum expected hit rate of 0.44 would have been observed (by design) in all conditions (i.e., 0.68 probability of the high-probability target multiplied by 65% target hits based on the performance-adjusted rewarded target width). Consistent with the Composite Benefit criterion, this is less than the expected hit rates observed experimentally in all conditions (Table 1, main diagonal; the evidence values for each row are 27.0, 74.7, 80.0, 61.6, and 13.7 dB). Subjects did not simply plan to reach to the most probable target but instead crafted a composite plan that allowed for the possibility that other, less-probable targets might be designated the reach target.

Scale experiment

In the Location experiment we found that subjects varied the spatial location of the point where the initial part of the reach crossed the trigger plane in response to changes in prior distributions, moving the fingertip closer to the peak of the prior probability distribution. Subjects deliberately traded off accuracy at hitting low-probability targets for improved accuracy at hitting high-probability targets. We could not reject the hypothesis that they chose optimal movement strategies for each prior (Row Dominance and Composite Benefit tests).

In the Scale experiment, we used three priors that shared the same central peak position but that differed in the width of the peak probability region (Fig. 2B). This set of priors varied the certainty with which the subject knew the location of the target before movement onset, while keeping the mean and median of the prior constant at the center of the distribution. Because

increasing the width of the prior increased the uncertainty of target location and therefore the probability of needing a trajectory adjustment to hit the target, we predicted that subjects would tend to decrease their speed at the trigger plane with increasing uncertainty of the prior, while maintaining a fingertip spatial trajectory similar to that observed when aiming toward the central target location during determinate-target reaches.

FINGERTIP SPATIAL TRAJECTORIES. In Fig. 4 we plot mean spatial trajectories by target for composite and simple reaches (across all subjects and conditions). Composite-reach spatial trajectories (closed circles) begin along the same trajectory found for simple reaches to the central target, both in their spatial coordinates (Fig. 4A) and in their direction (Fig. 4B). There is a leftward curvature during the main portion of the spatial trajectories, for both simple and composite reaches. This curvature is the result of a slow, approximately constant-magnitude change of movement direction throughout most of the reach, seen in Fig. 4B as the straight-line trajectory describing movement direction over the relevant portions of the reaches.

As described earlier, the instantaneous direction of fingertip motion toward each of the seven target positions during reaches to determinate targets is almost immediately distinct for distinct targets (Fig. 4). This differentiation is delayed in reaches to probabilistic targets for about 147–177 mm (corresponding to 150–196 ms after presentation of the target), depending on the criterion chosen.

Figure 5 plots the variance of the direction of fingertip motion (“directional variance”) at each position along the way to the screen. The filled black circles plot the directional variance pooled over all targets (over all data points at each y-position contributing to the average trajectories plotted as filled symbols in Fig. 4B). The open circles plot the variance calculated relative to the mean direction within each target condition (variance calculated over all differences between data points contributing to the filled symbols in Fig. 4B and the corresponding average trajectory direction for that target con-




FIG. 5. Onset of trajectory compensations in the Scale experiment. Fingertip motion angular variance of the test reaches (left ordinate) is plotted as a function of distance from the screen, calculated either relative to the mean of each condition and target (open circles) or relative to the overall mean across conditions and targets (filled circles). Evidence for a difference between these curves is overlaid and scaled to the right ordinate (filled diamonds). Two criteria for a significant difference are shown: evidence above zero (dotted line) and the start of the rise of evidence to its peak value (dashed line). Black vertical bar indicates the trigger plane.

⁵ Because diagonal elements can produce evidence that they are neither greater than nor less than themselves, diagonal evidence values must be 0 dB in all cases.

dition). The filled diamonds and right-hand ordinate indicate the evidence that these two variance values differ. At 196 ms after presentation of the target (mean distance of 177 mm, dotted line), the evidence function becomes positive. This is a reasonably conservative criterion for the onset of target differentiation in the movement given that we are looking for a pattern of results in which the evidence becomes greatest just before the target plane and decreases to a stable level before and after. A less-conservative estimate (150 ms, or 147 mm) results from a criterion based on the point at which the evidence function begins to rise to its peak value (Fig. 5, dashed line). Although the sign of the evidence calculated at that point is negative, the overall pattern argues that this is still a reasonable choice for the point of divergence toward individual targets. It is also worth mentioning that fingertip motion direction is a more sensitive measure of the initial divergence toward the final reach target than is the same analysis performed on horizontal spatial-position data. For example, using the criterion that the evidence function crosses zero as the start of divergence, the estimated latency based on position variance is 231 ms, 35 ms later than the estimate based on the same criterion derived from directional variance.

VELOCITY PROFILES. Forward velocity profiles peak shortly after the trigger plane is crossed, just before the halfway point of the reach, consistent with previous studies (e.g., Konczak and Dichgans 1997; Morasso 1981). These profiles displayed the roughly parabolic shape generally observed during similar reaching movements (e.g., Milner and Ijaz 1990; Nakano et al. 1999; Todorov and Jordan 1998), with the expected deviations from this pattern occurring near the end of reaches requiring substantial terminal corrections. That is, composite-reach forward velocity was slightly reduced during the lateral excursions required for large trajectory adjustments near the ends of some reaches. Nevertheless, reaches were always smooth and velocity profiles observed during composite reaches had shapes virtually identical to velocity profiles observed during determinate-target reaches.

In addition to calculating average velocity profiles, we determined whether there were significant carryover effects from one reach to the next on subjects' trigger plane crossing speeds. In other words, we asked whether a crossing point speed slightly above or below average on a given reach would be followed by a correction on one or more of the immediately subsequent reaches. There were no significant autocorrelations of trigger-plane crossing speeds beyond lag 0, indicating that the speed at which subjects' fingertips crossed the trigger plane on a given trial was unaffected by the crossing speeds experienced in previous trials.

Velocity at the trigger plane varied as a function of the information available in each of the prior probability distributions for target location. Subjects modulated their speed at the trigger plane such that they moved fastest when the probability distribution was most informative and slowed as the information content decreased (Fig. 6).

As with the results of the Location experiment, it is possible to generate and test the predictions of a "mixtures-of-strategies" model to the present results. Here, the prediction for a mixture of strategies is even more forcefully rejected than earlier because determinate-reach velocities at the trigger plane were all indistinguishable (all *P*-values were >0.1). If subjects had probabilistic-

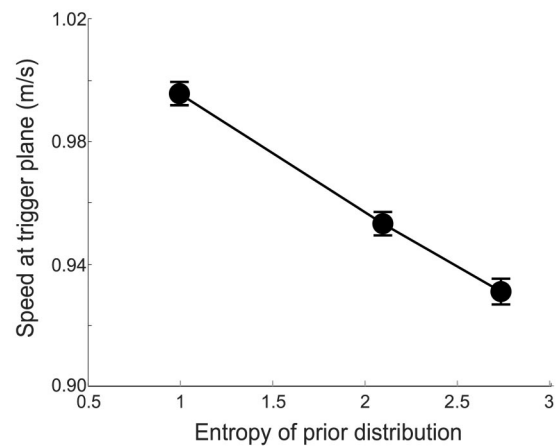


FIG. 6. Speed and information content. Speed at the trigger plane is plotted as a function of the entropy of the prior distribution for the Scale experiment. Error bars are SEs across subjects and targets for each prior probability distribution.

cally mixed the determinate-reach trajectories to produce their reach profiles in the three conditions examined here, there would have been no variation in trigger-plane crossing speed. This is in sharp contrast to the result shown in Fig. 6.

ACCELERATION PROFILES. Acceleration profiles were approximately linear during the main portion of the reaches, as would be expected from bell-shaped velocity profiles, excluding an initial sharp increase and a spike near the end of the profile as trajectories were adjusted near the target location. There were no strong differences between determinate and probabilistic reaches in the acceleration profiles or between acceleration profiles observed under the three prior probability distributions.

ROW DOMINANCE TEST. We next tested whether the adjustments made to the three levels of target certainty passed the Row Dominance test. As can be seen in Table 2 (see Supplement for confidence intervals surrounding estimates of expected hit rates), there is a failure of Row Dominance in the high-certainty condition. Our evidence analysis confirms this, indicating that performance using the observed strategy in the high-certainty condition produced significantly poorer performance than what would have been obtained from using the strategy used in the medium-certainty condition (strategy 2).

COMPOSITE BENEFIT TEST. In addition to being significantly suboptimal by the Row Dominance test, reach planning was also suboptimal by the Composite Benefit test (evidence values for each row of -116.3 , 460.5 , and 728.1 dB). A simple movement plan would have produced maximum hit rates of 0.72, 0.26, and 0.14 under the high-, medium-, and low-certainty conditions, respectively (these represent the product of the probability of one of the high-probability targets times the 85% hit rate based on the performance-adjusted rewarded target width). Subjects would therefore have obtained greater earnings had they used a simple movement plan in the high-certainty condition. In fact, the simple movement plan would have outperformed not only the observed composite trajectories in that condition, but also the best of the observed movement plans (s_2 , observed in the medium-certainty condition) in the high-certainty condition (see Table 2). By both the Row Dominance and Composite Benefit criteria, subjects' performance in the Scale experiment was suboptimal.

TABLE 2. *Scale Experiment: Row Dominance (pooled over subjects)*

	Strategy 1	Strategy 2	Strategy 3
Condition 1	0.58 (0)	0.62 (-13.4)	0.57 (3.3)
Condition 2	0.49 (19.8)	0.55 (0)	0.54 (2.5)
Condition 3	0.37 (61.4)	0.48 (-0.7)	0.48 (0)

Results of the Row Dominance test in the Scale experiment. Each element of the table is the probability of obtaining a hit when combining one of the three observed reach strategies (indexed by column) with one of the three prior distributions (indexed by row). The calculation of probabilities, evidence, and confidence intervals, as well as the organization, follow that of Table 1.

DISCUSSION

In the speeded reaching task considered here, the subject does not know the actual target of the movement until the fingertip has arrived at an invisible “trigger plane” approximately one third of the way between the starting point and the target. The subject does know the possible continuations to each possible target and the prior probability that each target will be the actual target. The challenge is to plan a mid-reach state specifying the location, velocity, and so forth of the fingertip at the trigger plane that is a compromise between the possible targets and that maximizes expected gain. The subject could plan a trajectory to the trigger plane that arrives at a particular location with a particular velocity and may plan higher derivatives of the trajectory as well.

We have presented experimental evidence that movement plans change in response to manipulation of prior probability distributions, and these changing movement plans serve to alter the location and speed of the mid-reach state of the arm. When we moved the location of the high-probability target in the Location experiment, subjects responded by planning trajectories to the trigger plane that differed primarily in location. In the Scale experiment, we increased the width of the high-probability center of the prior distribution (and thereby the uncertainty about true target location). In response, subjects reduced the speed of trajectories at the trigger plane.

We formulated and tested two criteria for optimal performance maximizing expected gain. We could not reject the hypothesis of optimal movement planning in the first (Location) experiment by either criterion. However, we could reject the hypothesis of optimal movement planning in the second (Scale) experiment. Subjects altered their trajectories in response to changes in target uncertainty, but not optimally.

Probabilistic anticipatory control and optimality

In the Location experiment, participants planned a movement that was a compromise between moving directly to the highest-probability target and moving to the central target. When new target information was provided after the reach passed the trigger plane, this led to an abrupt change in direction requiring an increase in torque. Lower torques generate lower levels of multiplicative motor noise (Hamilton and Wolpert 2002; Jones et al. 2002), ultimately leading to higher hit rates and greater expected gain. Thus one can interpret the results of our experiments in terms of the biomechanical constraints on good performance.

In either of the two experiments, we can imagine the continuation trajectories from any mid-state at the trigger plane to

any possible target and compare them according to the torque incurred in changing direction. If the subject plans to move to the trigger plane at the far left edge then continuation trajectories that return to targets at the right-hand side will involve a large change in direction of travel. A change in location at the trigger plane in turn changes the torque-induced movement error and ultimately the probability of hitting a possible target on trials when it proves to be the actual target. Moreover, the faster the fingertip is moving at the target plane, the greater torque incurred in changing direction. The ideal movement planner must choose a movement plan to trade off torque-induced movement error for continuations to high- and low-probability targets.

An implication of our results is that subjects are able to implement a predictive control strategy that takes into account the probability of later trajectory changes, integrating early probabilistic target information with knowledge of biomechanical and neural constraints. These findings are consistent with other recent work demonstrating compensatory torques for Coriolis and other anticipated perturbing forces (Flanagan and Wing 1997; Hudson et al. 2005; Kim et al. 2006; Lackner and DiZio 1994; Patla et al. 2002; Pigeon et al. 2003a; Scheidt et al. 2005; Tunik et al. 2003; Wang and Sainburg 2005). The current results show, in addition, that these adaptive precompensations can influence the velocity profile of the reach and not just the spatial trajectory or endpoint of the reach.

Previous work suggesting predictive compensatory torque generation assumed a deterministic computation of the magnitude of compensation based on the physics of the to-be-compensated forces. Although it is clear that a predictable contingency must be present for the planning of compensatory torques, this does not necessarily imply an internal model based simply on a deterministic physical relationship. Our results show that predictive control is also influenced by both the probabilities that compensatory torques will be required and the expected magnitudes of those torques.

Suboptimality of speed modulation

Given that subjects did not modulate speed in a manner consistent with that of an optimal movement planner, it is interesting to speculate about the possible causes of this suboptimality. Because the pattern of speed modulation was exactly as predicted, our main clues regarding the suboptimality are provided by the Composite Benefit and Row Dominance tests. The Composite Benefit test tells us that, in the high-certainty condition, subjects would have improved their performance by ignoring all but the central target, and simply concentrating on hitting that target whenever it appeared. That is, subjects appear to place too high a value on hitting the occasional eccentric target. In addition, Table 2 tells us that performance would have increased in this condition by reducing speed at the trigger plane—speed modulation with target uncertainty was greater than what would have been required for subjects to maximize gain.

It is possible that the suboptimality observed here is due to a lack of fidelity in subjects' representations of the prior probability distribution of target locations. This representation was learned by experiencing each probability distribution of target locations during the determinate-target reaches made at the beginning of each subject's session. Although this implicit learning may have led to

an imperfect representation of the relevant prior probability distributions, it is unclear why this would occur only in the Scale but not in the Location experiment.

Timing of deviations toward noncentral targets

We detected subjects' responses to target presentation at latencies of 150–195 ms (direction) or 171–231 ms (position), depending on the criterion chosen (Fig. 5). These latencies are consistent with estimates of simple reaction times (150–200 ms) and those measured by Soechting and Lacquaniti (1983) in a two-step paradigm (who found latencies of 150–200 ms measured by the initial EMG response and 180–230 ms for a change in kinematic variables). This is perhaps surprising, given the greater complexity of our task, in which subjects were required to detect the target and then compute and put into effect the torques needed for the requisite change in trajectory. However, most studies have used a detection criterion based on position, which we found to be less sensitive than our movement-direction-based criterion (Fig. 5).

Characteristic reach paths

In our study, as in previous research, we found slightly curved reach trajectories (Figs. 3 and 4). This finding relates to the issue of directional versus positional control of reaching, and the oft-cited description of reach trajectories as following a “straight-line path.” This is an inaccurate description of normal reach trajectories, as previously noted (Goodbody and Wolpert 1999). It is an open question whether the curvature of normal reaches projected onto a horizontal plane is due to some aspect of the perception of the movement (Brenner et al. 2002; Flanagan and Rao 1995; Goodbody and Wolpert 1999; Osu et al. 1997) or to intrinsic biomechanical (Goodbody and Wolpert 1999) or computational (Osu et al. 1997) factors. In this respect, it is of interest to note that although the trajectory is curved, in our data movement direction is a linear function of the distance traveled. In the context of spatial trajectories, it has been argued that planning takes place in the coordinate system in which the description of the movement is a straight line (e.g., Nakano et al. 1999; Pigeon et al. 2003b). This line of argument would lead to the conclusion that it may be direction and not position that is planned in our reaching task.

Planning single- and multiple-target movements

When planning movements to a series of targets, it is possible that each segment of the movement is planned independently of the other segments. This would be an optimal strategy for a reach to N targets in succession if it were impossible to adjust the state of the fingertip for any of the i ($i < N$) intermediate targets to increase the probability of acquiring the $(i + 1)$ target. In a study of speeded reaching to two consecutive targets visible before reach initiation, Aivar et al. (2005) found that the reach segments to the two targets were not planned independently.

In relation to the current study, we note that the trigger plane is similar to an initial target, but one that it is impossible to avoid on the way to the second target (the screen), and for which the state of the fingertip as it is reached can be planned with many fewer constraints than would be possible using an initial target that was spatially constrained. If the constraints on

the state of the fingertip as it passed through the trigger plane were made more restrictive, the movement plan governing each segment of the reach could be formed more independently of the plans for other segments. That is, the movement plan for the entire sequence would tend to resemble a series of simple-movement plans instead of a composite movement plan. There is no advantage to forming a composite movement plan if there is no way to bias one state of the fingertip to facilitate achieving another state.

Conclusions

We investigated how human subjects plan speeded reaching movements when the exact target of the reach is not known during the initial part of the movement. At the start of each trial, subjects see an array of potential targets (vertical bars) for the reaching movement. Any of the targets could be the actual target for that trial and, initially, the subject is given only the prior probability that each potential target could be the actual target. After the subject has moved one third of the distance to the screen (and his/her fingertip has passed through an invisible “trigger plane”) the actual target is marked. If the subject touches the actual target on a trial within the time limit, he or she earns a monetary reward.

The challenge for the subject is to plan the initial part of the movement to the trigger plane without knowing the location of the actual target. This initial movement has many possible continuations, to each of the possible targets. The subject knows the prior probability that each possible continuation will lead to the actual target and must select a composite movement that strikes a balance between possible continuations. As the prior distribution changes, the subject may alter the location, velocity, and possibly higher derivatives of fingertip location at the trigger plane.

We manipulated the prior distribution of potential targets in two experiments and measured how location and speed of the fingertip changed at the trigger plane. In the Location experiment, one potential target was more probable than the remaining potential targets, all of which were equally likely. We manipulated the location of the most probable target and examined how subjects varied the mean location and velocity at which they passed through the trigger plane. We expected that subjects would primarily alter location but not velocity and that is what we found. In the Scale experiment, the prior consisted of a central higher-probability region and symmetric, surrounding lower-probability regions. We varied the width of the central part of the prior distribution, thereby increasing or decreasing the uncertainty (entropy). We found that increasing uncertainty led subjects to arrive at the trigger plane at lower velocities.

For our purposes, an “ideal” or “optimal” movement planner is an algorithm that plans movements to maximize expected gain. In our task, an ideal movement planner would plan a composite movement that places the fingertip in the trigger plane at a location and traveling at a speed that represents the trade-off between possible continuations of the movement and their probabilities that maximizes expected gain.

We developed two necessary conditions that an ideal movement planner must satisfy and tested whether subjects satisfied them. The first was the Row Dominance criterion. We computed how well subjects would have done in each condition if they had adopted the strategy that they used in each of the other conditions.

The optimal movement planner, by definition, picks the optimal strategy in each condition. Consequently, if we find that subjects in any condition could have earned more on average by adopting the strategy they used in another condition, we can reject the hypothesis that they are optimal. Although this was not the case in the Location experiment (we found no evidence that subjects could have improved their performance by choosing another of the observed strategies), a clear pattern of deviation from optimality was observed in the Scale experiment.

Suboptimal performance was also detected in the Scale experiment by the Composite Benefit test. The results of this test mirrored the results of the Row Dominance test: in the Scale experiment, this test failed as well, but we found no evidence for failure in the Location experiment.

We have therefore demonstrated that subjects plan position and velocity at an arbitrary mid-reach location based on probabilistic information provided before reach initiation. We could not reject the hypothesis of optimality by both the Composite Benefit and the Row Dominance tests in the Location experiment. However, both tests reject the hypothesis that subjects optimally planned velocity at the trigger plane in the Scale experiment.

GRANTS

This work was supported by National Eye Institute Grant EY-08266.

REFERENCES

- Abrams RA, Meyer DE, Kornblum S.** Eye-hand coordination: oculomotor control in rapid aimed limb movements. *J Exp Psychol Hum Percept Perform* 16: 248–267, 1990.
- Aivar MP, Brenner E, Smeets JBJ.** Correcting slightly less simple movements. *Psicológica* 26: 61–79, 2005.
- Bédard P, Proteau L.** On-line vs. off-line utilization of peripheral visual afferent information to ensure spatial accuracy of goal-directed movements. *Exp Brain Res* 158: 75–85, 2004.
- Brainard DH.** The psychophysics toolbox. *Spat Vis* 10: 433–436, 1997.
- Brenner E, Smeets JB.** Colour vision can contribute to fast corrections of arm movements. *Exp Brain Res* 158: 302–307, 2004.
- Brenner E, Smeets JB, Remijnse-Tamerius HC.** Curvature in hand movements as a result of visual misjudgments of direction. *Spat Vis* 15: 393–414, 2002.
- Elliott D, Binsted G, Heath M.** The control of goal-directed limb movements: correcting errors in the trajectory. *Hum Mov Sci* 18: 121–136, 1999.
- Flanagan JR, Rao AK.** Trajectory adaptation to a nonlinear visuomotor transformation: evidence of motion planning in visually perceived space. *J Neurophysiol* 74: 2174–2178, 1995.
- Flanagan JR, Wing AM.** The role of internal models in motion planning and control: evidence from grip force adjustments during movements of hand-held loads. *J Neurosci* 17: 1519–1528, 1997.
- Goodbody SJ, Wolpert DM.** The effect of visuomotor displacements on arm movement paths. *Exp Brain Res* 127: 213–223, 1999.
- Gribble PL, Mullin LI, Cothros N, Mattar A.** Role of cocontraction in arm movement accuracy. *J Neurophysiol* 89: 2396–2405, 2003.
- Hamilton AF, Jones KE, Wolpert DM.** The scaling of motor noise with muscle strength and motor unit number in humans. *Exp Brain Res* 157: 417–430, 2004.
- Hamilton AFC, Wolpert DM.** Controlling the statistics of action: obstacle avoidance. *J Neurophysiol* 87: 2434–2440, 2002.
- Haruno M, Wolpert DM, Kawato M.** MOSAIC model for sensorimotor learning and control. *Neural Comput* 13: 2201–2220, 2001.
- Heath M, Westwood DA, Binsted G.** The control of memory-guided reaching movements in peripersonal space. *Motor Control* 8: 76–106, 2004.
- Hudson TE, DiZio P, Lackner JR.** Rapid adaptation of torso pointing movements to perturbations of the base of support. *Exp Brain Res* 165: 283–293, 2005.
- Jaynes ET.** *Probability Theory: The Logic of Science*. Cambridge, UK: Cambridge Univ. Press, 2003.
- Jeffreys H.** An invariant form for the prior probability in estimation problems. *Proc R Soc Lond A Math Phys Eng Sci* 186: 453–461, 1946.
- Jones KE, Hamilton AF, Wolpert DM.** Sources of signal-dependent noise during isometric force production. *J Neurophysiol* 88: 1533–1544, 2002.
- Kim SW, Shim JK, Zatsiorsky VM, Latash ML.** Anticipatory adjustments of multi-finger synergies in preparation for self triggered perturbations. *Exp Brain Res* 174: 604–612, 2006.
- Komilis E, Péllisson D, Prablanc C.** Error processing in pointing at randomly feedback-induced double step stimuli. *J Mot Behav* 25: 299–308, 1993.
- Konczak J, Dichgans J.** The development toward stereotypic arm kinematics during reaching in the first 3 years of life. *Exp Brain Res* 117: 346–354, 1997.
- Körding KP, Wolpert DM.** Bayesian integration in sensorimotor learning. *Nature* 427: 244–247, 2004.
- Lackner JR, DiZio P.** Rapid adaptation to Coriolis force perturbations of arm trajectory. *J Neurophysiol* 72: 299–313, 1994.
- Milner TE, Ijz MM.** The effect of accuracy constraints on three-dimensional movement kinematics. *Neuroscience* 35: 365–374, 1990.
- Morasso P.** Spatial control of arm movements. *Exp Brain Res* 42: 223–227, 1981.
- Nakano E, Imamizu H, Osu R, Uno Y, Gomi H, Yoshioka T, Kawato M.** Quantitative examinations of internal representations for arm trajectory planning: minimum commanded torque change model. *J Neurophysiol* 81: 2140–2155, 1999.
- Osu R, Uno Y, Koike Y, Kawato M.** Possible explanations for trajectory curvature in multijoint arm movements. *J Exp Psychol Hum Percept Perform* 23: 890–913, 1997.
- Patla AE, Ishac MG, Winter DA.** Anticipatory control of center of mass and joint stability during voluntary arm movement from a standing posture: interplay between active and passive control. *Exp Brain Res* 143: 318–327, 2002.
- Péllisson D, Prablanc C, Goodale MA, Jeannerod M.** Visual control of reaching movements without vision of the limb. II. Evidence of fast unconscious processes correcting the trajectory of the hand to the final position of the double step stimulus. *Exp Brain Res* 62: 303–311, 1986.
- Pelli DG.** The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat Vis* 10: 437–442, 1997.
- Pigeon P, Bortolami SB, DiZio P, Lackner JR.** Coordinated turn-and-reach movements. I. Anticipatory compensation for self-generated Coriolis and interaction torques. *J Neurophysiol* 89: 276–289, 2003a.
- Pigeon P, Bortolami SB, DiZio P, Lackner JR.** Coordinated turn-and-reach movements. II. Planning in an external frame of reference. *J Neurophysiol* 89: 290–303, 2003b.
- Rabin E, Gordon AM.** Influence of fingertip contact on illusory arm movements. *J Appl Physiol* 96: 1555–1560, 2004.
- Sabes PN, Jordan MI.** Obstacle avoidance and a perturbation sensitivity model for motor planning. *J Neurosci* 17: 7119–7128, 1997.
- Saunders JA, Knill DC.** Visual feedback control of hand movements. *J Neurosci* 24: 3223–3243, 2004.
- Scheidt RA, Condit MA, Secco EL, Mussa-Ivaldi FA.** Interaction of visual and proprioceptive feedback during adaptation of human reaching movements. *J Neurophysiol* 93: 3200–3213, 2005.
- Schmidt T.** The finger in flight. Real time motor control by visually masked color stimuli. *Psychol Sci* 13: 112–118, 2002.
- Soechting JF, Lacquaniti F.** Modification of trajectory of a pointing movement in response to a change in target location. *J Neurophysiol* 49: 548–564, 1983.
- Todorov E.** Cosine tuning minimizes motor errors. *Neural Comput* 14: 1233–1260, 2002.
- Todorov E, Jordan MI.** Smoothness maximization along a predefined path accurately predicts the speed profiles of complex arm movements. *J Neurophysiol* 80: 696–714, 1998.
- Todorov E, Jordan MI.** Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5: 1226–1235, 2002.
- Torres EB, Zipser D.** Simultaneous control of hand displacements and rotations in orientation-matching experiments. *J Appl Physiol* 96: 1978–1987, 2004.
- Trommershäuser J, Maloney LT, Landy MS.** Statistical decision theory and trade-offs in the control of motor response. *Spat Vis* 16: 255–275, 2003a.
- Trommershäuser J, Maloney LT, Landy MS.** Statistical decision theory and the selection of rapid, goal-directed movements. *J Opt Soc Am A* 20: 1419–1433, 2003b.
- Tunik E, Poizner H, Levin MF, Adamovich SV, Messier J, Lamarre Y, Feldman AG.** Arm-trunk coordination in the absence of proprioception. *Exp Brain Res* 153: 343–355, 2003.
- Vindras P, Viviani P.** Altering the visuomotor gain. Evidence that motor plans deal with vector quantities. *Exp Brain Res* 147: 280–295, 2002.
- Wang J, Sainburg RL.** Adaptation to visuomotor rotations remaps movement vectors, not final positions. *J Neurosci* 25: 4024–4030, 2005.
- Woodworth RS.** The accuracy of voluntary movement. *Psychol Rev Monogr* 3: 1–114, 1899.